NAVAL POSTGRADUATE SCHOOL
Monterey, California

THESIS

**PROTECTING THE FORCE:
APPLICATION OF STATISTICAL PROCESS CONTROL
FOR FORCE PROTECTION IN BOSNIA**

by

Paul J. Finken

June 2000

Thesis Advisor:	David H. Olwell
Second Reader:	Thomas W. Lucas

Approved for public release; distribution is unlimited

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	June 2000	Master's Thesis

4. TITLE AND SUBTITLE
Protecting the Force: Application of Statistical Process Control for Force Protection in Bosnia

5. FUNDING NUMBERS

6. AUTHOR(S) : Paul J. Finken

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Naval Postgraduate School
Monterey, CA 93943-5000

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSORING / MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES
The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

12a. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT *(maximum 200 words)*

Military commanders determine the appropriate Force Protection measures to protect their units from a wide variety of threats based on their assessment of the enemy threat in the specific situation. They currently have no statistical tool from which to base their assessment of the threat, or to recognize changes in the current situation. In Operations Other Than War (OOTW), environments where the enemy is disorganized and incapable of mounting a deception plan, staffs could model hostile events as stochastic events and use statistical methods to detect changes to the process. This thesis developed a statistical tool, based on Cumulative Sum (CUSUM) and Shewhart Charts, that military leaders can use in OOTW environments to recognize statistically significant changes in the situation. The tool applies current univariate control chart methods, as well as an original nonparametric multivariate control scheme developed in this thesis, to North Atlantic Treaty Organization (NATO) Stabilization Force (SFOR) incident data. The tool enables commanders to identify isolated and persistent shifts in the means of the data categories or shifts in the correlation of three data categories. By recognizing changes in the current situation, military leaders have a basis from which to change their force protection measures and better protect their unit.

14. SUBJECT TERMS
Statistical Process Control, Multivariate, Force Protection

15. NUMBER OF PAGES 124

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFI- CATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	UL

NSN 7540-01-280-5500

Standard Form 298 (Re2-89)
Prescribed by ANSI Std. 239-18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

PROTECTING THE FORCE:
APPLICATION OF STATISTICAL PROCESS CONTROL
FOR FORCE PROTECTION IN BOSNIA

Paul J. Finken
Major, United States Army
B.S., United States Military Academy, 1989

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the

NAVAL POSTGRADUATE SCHOOL
June 2000

Author: ______________________
Paul J. Finken

Approved by: ______________________
David H. Olwell, Advisor

____
Thomas W. Lucas, Second Reader

____
Richard E. Rosenthal, Chairman
Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Military commanders determine the appropriate Force Protection measures to protect their units from a wide variety of threats based on their assessment of the enemy threat in the specific situation. They currently have no statistical tool from which to base their assessment of the threat, or to recognize changes in the current situation. In Operations Other Than War (OOTW), environments where the enemy is disorganized and incapable of mounting a deception plan, staffs could model hostile events as stochastic events and use statistical methods to detect changes to the process. This thesis developed a statistical tool, based on Cumulative Sum (CUSUM) and Shewhart Charts, that military leaders can use in OOTW environments to recognize statistically significant changes in the situation. The tool applies current univariate control chart methods, as well as a new nonparametric multivariate control scheme developed in this thesis, to SFOR incident data. The tool enables commanders to identify isolated and persistent shifts in the means of the data categories or shifts in the correlation of three data categories. By recognizing changes in the current situation, military leaders have a basis from which to change their force protection measures and better protect their unit.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF FIGURES

LIST OF TABLES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS

ARL	Average Run Length
CALL	Center for Army Lessons Learned
CUSUM	Cumulative Sum
DI	Decision Interval
FIR	Fast Initial Response
IPB	Intelligence Preparation of the Battlefield
LCL	Lower Control Limit
NATO	North Atlantic Treaty Organization
OOTW	Operations Other Than War
SFOR	Stabilization Force
SPC	Statistical Process Control
S2	Battalion Intelligence Officer
UCL	Upper Control Limit

THIS PAGE INTENTIONALLY LEFT BLANK

EXECUTIVE SUMMARY

Tactical commanders in the Army rely on pattern recognition methods to detect changes to the current situation, which in turn form the basis for their tactical decisions and plans. Commanders do not have a tool that enables them to differentiate the naturally occurring random variations in the situation from statistically significant changes in the situation. In Operations Other Than War (OOTW), where the enemy is disorganized and incapable of mounting a deception plan, staffs could model hostile events as stochastic events and use statistical methods to detect significant changes in the situation.

This thesis, specifically targeted at units deployed to Bosnia as part of the North Atlantic Treaty Organization (NATO) Stabilization Force (SFOR), developed a statistical tool that allows military leaders to analyze enemy incident data and determine when statistically significant changes in the situation occur. The tool is implemented in an Excel worksheet with Visual Basic macros, and is based on statistical process control (SPC) Cumulative Sum (CUSUM) and Shewhart control charts. The tool's graphical and text outputs ensure easy identification of the shifts and the time periods in which they occur.

The methods used in the worksheet utilize current SPC techniques for analyzing univariate Poisson data and also a nonparametric method for analyzing multivariate data, developed in this thesis. The univariate Poisson methods enable commanders to analyze predictor variables separately to detect isolated departures and persistent shifts in the mean number of the individual variables. The nonparametric multivariate method enables them to analyze three predictor variables simultaneously to detect isolated departures and persistent shifts in the mean number of predictor variables, as well as isolated departures and persistent shifts in the correlation structure of the variables.

In the case of the SFOR in Bosnia, actions of the different ethnic

groups from March to October 1999 are tabulated and categorized into three categories: threats and rhetoric, contentious activities, and violent actions toward SFOR. We analyzed the data using the methods described above to identify statistically significant isolated departures and statistically significant persistent shifts in the data categories. By identifying statistically significant changes in the situation, the commander is able to make more informed decisions and appropriate changes to the force protection level of his unit.

Results from the analysis suggest several key issues about the situation that the commander should find informative and useful when developing his force protection plan. First, the situation was the most hostile in the initial data collection periods, 1 March through 5 April 1999, as denoted by high number of incidents in all data categories. The high numbers of enemy incidents were not naturally occurring random variations in the situation, but were instead statistically significant isolated departures from the usually observed values. In particular, statistically significant high numbers of incidents occurred in category 3, violence towards SFOR, from 22 through 28 March, and in category 3, threats and rhetoric, from 29 March through 4 April. Possible causes for these increases may be found in the fact that they coincide with the United Nation's efforts to broker a peace settlement in Kosovo from February through the middle of March 1999, and the NATO air strikes against Serbian facilities, which commenced on 25 March 1999. Looking at the SFOR incident log during 22 through 28 March, which corresponds to the start of the bombing campaign, reveals that at least six of the eleven demonstrations against SFOR were anti-bombing demonstrations. From 29 March through 4 April, the number increased to 12 out of 17.

The high levels of enemy incidents explained above were isolated occurrences, with the numbers of incidents decreasing rapidly after 5 April. Increasing force protection levels after these incidents occurred would be somewhat ineffective. The changes would not take

effect until after the highest threat has already passed. Increasing force protection level will be effective in protecting the force against the lesser threats that occur as the number of incidents decrease.

Commanders should not be completely convinced by this seemingly obvious cause of the high number of incidents. They should proceed with additional analysis of the situation to determine if other factors were present that may have caused or assisted in the increased number of incidents. The commander should use these factors to predict future enemy threat levels in similar situations. From these predictions, he can initiate the appropriate force protection levels prior to the situation occurring, thus better protecting his unit.

The initial high hostility period was followed by a continual decrease in the number of enemy incidents in all data categories through the end of the data collection period, 3 October 1999. The number of incidents decreased rapidly from 5 through 24 April. After 25 April, the numbers of incidents appeared to stabilize. The tool developed in this thesis however, identified numerous statistically significant persistent decreases in the number of incidents after 25 April. Two statistically significant decreases occurred in category 1, threats and rhetoric, and one statistically significant decrease occurred in each of category 2, contentious activities, and category 3, violence towards SFOR. All of these persistent decreases justify consideration of lower force protection levels of the unit. The commanders and their staffs need to analyze the situation further to determine the specific causes of these decreases and the appropriate force protection levels. By identifying the possible causes of these decreases, commanders could also focus their peacekeeping efforts in order to continue these trends.

It should be noted that there was an isolated statistically significant increase in the number of incidents in category 1, threats and rhetoric, from 13 through 19 September. As with other isolated increases discussed earlier, the cause of this increase should be

determined and used for future reference.

Finally, the correlation between the data categories did not change. That is to say, the enemy's efforts, as divided among the three categories, remained constant. This can be seen by the simultaneous increasing or decreasing trends that occurred in all three data categories. If a change in the correlation between the data categories was detected, it would indicated a change in the enemy's distribution of effort, say from threats to acts of violence. This information would be vital to the commander in his assessment of the threat and his determination of appropriate force protection levels.

Overall recommendations after analyzing the SFOR incident data are that the force protection measures be reduced due to the statistically significant decreases in the number of enemy incidents after 5 April 1999. However, sufficient protection should be maintained to safeguard against possible isolated increases in enemy incidents, as detected in category 1, threats and rhetoric, 13 through 19 September.

As shown above, the tool developed in this thesis provides vital information about the enemy situation that may not have otherwise been obtainable by the commander. It enables the commander to quickly differentiate between normal random variation in the situation and statistically significant changes in the situation. This will greatly assist the commander in assessing the enemy threat and developing his force protection plan. This tool is not an omniscient tool by which commanders can guarantee the 100% safety of their soldiers. It is, however, the first and only statistical tool that the commander has at his disposal for detecting changes in the enemy situation.

ACKNOWLEDGEMENT

This thesis would not be complete without the acknowledgement of all those that assisted me in this endeavor. Without their help, this work would not have been possible.

First and foremost, I thank my thesis advisor, LTC Olwell, for mentoring and teaching me in every aspect of this thesis. His insight and expertise in the field of statistics and Statistical Process Control was vital to this work. I thank my second reader, Professor Lucas, whose thorough review ensured my success. I also thank LTC Kastner and LTC Itao, as well as the rest of the Task Force Eagle staff, who provided the data that made the analysis possible. A special thanks goes out to LT Mitch Kerman whose expertise in Visual Basic programming helped me clear some of the programming hurdles I encountered in a timely manner. I thank my parents for giving me the confidence and determination that enables me to accomplish anything. And finally I thank my wife Jackie and daughter Emilie for motivating and supporting me through the many months of this work.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

Force protection is defined as the "security plan designed to protect soldiers, civilian employees, family members, facilities and equipment in all locations and situations…" (Department of the Army, 1994, p106). Its primary focus is to sustain the strength of the force in order to accomplish the mission. It is a key planning consideration in all operations from high intensity conflict to daily soldier training, and should consider every possible threat from terrorist attacks to simple disease prevention.

In conventional combat operations, the enemy is organized and conducts operations in accordance with its doctrine. This normally includes the use of deception, displaying a false posture, to assist in ensuring the success of the main effort. The friendly commander uses the Intelligence Preparation of the Battlefield (IPB) process to assess the enemy capabilities and determine how best to defeat him. In the IPB process, the friendly commander gathers intelligence to determine the enemy's position, strength, and capabilities. He then compares this to the enemy's doctrine to predict the enemy's next course of action, to include when and where it will occur. (Department of the Army, 1990, p4-3) Facing an organized enemy, the commander must consider the enemy's use of deception throughout the entire IPB process. He cannot view the information collected as an absolute indicator of what the enemy is planning to do next. Since all actions for both the enemy and friendly are planned using strategy and a partial amount of information on the other side, game theory methods are best suited to model the actions of the opposing sides in this situation.

In Operations Other Than War (OOTW), however, the enemy consists of "loosely organized groups of irregulars, terrorists, or other conflicting segments of a population as predominate forces" (Department of the Army, 1994, pV). These loosely organized groups have no

predetermined doctrine (Department of the Army, 1993, p3-2), and in most cases their minimal command structures are incapable of coordinating a sophisticated deception plan. In the absence of doctrine, the friendly commander must create models based on enemy operational patterns. He develops operational patterns on the enemy by determining a set of events, or indicators, that best capture the character or operating habits of the enemy. He then establishes a record of these events by time and location, and analyzes these records to identify patterns in the events (Center for Army Lessons Learned, 1996, pp1-2). The commander and his staff use these patterns to predict future enemy events. Because the enemy is assumed to be incapable of executing a deception plan, the commander can view and model the events collected as tangible, stochastic indicators of future enemy actions. Because the events are stochastic, statistical methods are well suited to analyze and model this situation.

Unfortunately, commanders and their staff do not possess a statistical tool to determine if a change in the frequency of one of the indicators constitutes a statistically significant change in the situation. That is, if the change is the result of an actual shift in the frequency or is the result of normal stochastic variation in the situation. Such a tool would assist them in maximizing the speed of detection of these changes and in minimizing the occurrence of false alarms, i.e. thinking that a change had occurred when in fact it did not. This in turn will provide the commanders an opportunity to prudently adjust their force protection measures.

II. BACKGROUND

The catastrophic results of improper force protection measures are evident in the June 25, 1996 bombing of the U.S. Air Force Khobar Tower housing complex in Saudi Arabia, where 19 American service members were killed. In this incident, earlier terrorist activities, namely a car bomb in November 1995, signaled a possible increase in the terrorist threat targeted against U.S. forces. As a result, the U.S. Commander in Chief for the Central Command declared a "high" threat level for the entire country. Upon notification of the increased threat level, commands across Saudi Arabia initiated vulnerability assessments on all installations to include Khobar Towers. From these assessments, numerous force protection improvements were made. However, an investigation following the disaster concluded that even with all this information, the staff did not provide proper guidance to the commander of the unit, and that the commander failed to adequately protect his forces (Cohen, 1997, pp1-3).

As a result of the tragedy at Khobar towers, the Secretary of Defense, William J. Perry, issued a memorandum to the Chairman of the Joint Chiefs of Staff that stated, "this incident and others that almost certainly will follow demand an increased emphasis on force protection throughout the Department of Defense" (Perry, 1996, p1). From this new emphasis, local commanders were given increased responsibility and authority for force protection (Air Force News, 1996, p2) and new intensified training requirements were established for all deploying personnel.

Lessons learned in training exercises for units deploying to Bosnia have identified that although "S2s generally have a system for plotting incident overlays" they do not have a method of collating and analyzing the information to determine increasing threats or to develop threat models. The lessons learned also state that a "simple computer

database program can be used to more quickly discern patterns" (Center for Army Lessons Learned, 1996, p1). The Center For Army Lessons Learned (CALL) advises the S2 to enter the information into the computer on a series of fields and "use the computer to determine correlations between events and within a type of event" (Center for Army Lessons Learned, 1996, p1). Even though these points have been identified, no model or computer package has been constructed assist commanders in identifying the enemy threat and making the necessary force protection changes.

III. PURPOSE AND RATIONALE

In Bosnia and other OOTW environments, commanders can capitalize on the enemy's lack of deception by monitoring hostile events as stochastic indicators of the current situation. A statistical model that monitors and detects changes to the situation, both increases and decreases in the number and type of enemy incidents, would give the commander a tangible warning of a change in the situation and an opportunity to review his force protection measures. As stated above, the need for such a model exists and this need will become more pressing as the number of OOTW missions increases.

By monitoring numerous indicators ranging from small gestures to significant violent activities, commanders in Bosnia can get a complete picture of the threat they face. The incidents of small gestures, which are likely to occur often and may be overlooked by the commander, may serve as a predictor for the likelihood of an occurrence of an act of considerable violence, such as an outright attack against a SFOR base that resembles the Khobar Towers bombing.

Such a predictive model would be extremely useful in Bosnia and would fill a void in the SFOR's IPB and force protection assessment processes. It would allow commanders to monitor those indicators that are important at their specific level. It would prove extremely useful to units in Bosnia who are dealing with three separate warring factions who are undistinguishable from each other and are intermingled throughout the local populace.

THIS PAGE INTENTIONALLY LEFT BLANK

IV. METHODOLOGY

A. BASIC UNIVARIATE CONTROL CHART METHODS

1. Basic Control Chart Methods

Control Charts are used extensively throughout industry to monitor production processes to identify instability and unusual circumstances (Devore, 1995, p685). They enable managers to distinguish between random fluctuations in the process and a change in the process mean or variance. Typical control charts plot the data X_i, or a function of the data $a(X_i)$, versus calculated upper and lower control limits (Weitzman, 1999, p7). If the plotted data stays between the control limits, the process is considered in statistical control. If a data plot extends outside these limits, then the process is considered out of statistical control and it signals that variation other than the usual amount is present in the process. Control charts enable managers to quickly identify when the process has gone out of control while preventing them from making unnecessary interventions in the process when it is in control. This is valuable because huge profits can lost by shutting down a production line for a week to retool suspected faulty equipment when the equipment is in fact functioning properly and the end product is within specifications. Of course, equipment and manufacturing processes will not run forever without repair. Control charts assist the managers in identifying when the repairs are needed. No single chart completely captures all possible shifts in the variability in a process, but Shewhart style control chart and cumulative sum (CUSUM) charts are two extensively used charts that offer different but extremely complementary information (Hawkins and Olwell, 1998, p71).

The Shewhart style control chart is very effective for detecting isolated special causes that lead to large shifts in the data (Hawkins and Olwell, 1998, p7). It does this by testing the mean of a specific characteristic of the product from batches of the product. Isolated or

transient shifts in a process are somewhat common and can occur from numerous sources within the process. For example, consider taking 20 samples of five bolts each and measuring the hardness of the five bolts. If one of the samples was produced from a contaminated shipment of iron ore that resulted in the bolts not meeting the required average hardness specifications, the mean hardness of the sample would be lower than the other 19. If the mean hardness of this sample is outside the range of usual variation around the true mean, the Shewhart chart will identify this difference by plotting the batch mean outside the control limits. If the subsequent sample is taken from bolts made from acceptably pure iron ore resulting in a mean average hardness close to the true mean, the Shewhart chart will show that the batch mean and the process are in control (Hawkins and Olwell, 1998, p7).

Shewhart charts have one major limitation in that they are ineffective in detecting moderate persistent shifts in the data (Hawkins and Olwell, 1998, p7-9). Returning to the bolt example, if over the life of the machinery the threading tool used to thread the bolts to the correct diameter becomes worn, the resulting bolt diameters may slowly increase. The slight change in average bolt diameters of a particular batch will not be significant enough to cause an isolated out of control signal on the Shewhart chart. Personnel specifically trained on Statistical Process Control (SPC) may be able to detect this small shift by viewing the Shewhart chart and identifying a trend, but the typical process manager will not. CUSUM charts are often used in conjunction with the Shewhart charts to offset this shortcoming because they are better suited to detected moderate persistent step shifts in process parameters (Hawkins and Olwell, 1998, p71).

CUSUM charts are "tuned" to monitor data from a specific distribution and to detect a shift in the process mean (Hawkins and Olwell, 1998, p138). As with Shewhart charts, CUSUM charts plot data and control limits against time. The data that CUSUM charts plot,

however, is a calculated cumulative statistic S_n, not the raw data as in Shewhart charts.

This thesis uses the decision interval form of the CUSUM. This form facilitates visual identification of shifts in process mean (Hawkins and Olwell, 1998, p24). The decision interval form of the CUSUM is defined by the recursion:

$$S_0^+ = 0$$
$$S_0^- = 0$$
$$S_n^+ = \max(0, S_{n-1}^+ + X_n - k^+)$$
$$S_n^- = \min(0, S_{n-1}^- + X_n + k^-)$$

(Hawkins and Olwell, 1998, p25-26)

where S^+ monitors upward shifts in the process mean, S^- monitors downward shifts in the process mean, X_n is the observation, μ is the process mean, and n is the current iteration number. The k's listed above are different and are commonly distinguished as k^+ for the upward shift and k^- for the downward shifts. As the equations are written above, k^+ is a positive reference value and k^- is a negative reference value. Some care should be taken, as certain users prefer to use non-negative values of k's in their calculations. In this case, k^- is subtracted instead of added.

If the process follows a given distribution with a constant mean and standard deviation, the values of S_n can be considered a random walk with reflection at the horizontal axis. A line formed by the plotted S_n's will have an expected cumulative slope of 0 and will infrequently go outside the control limits. Once the process mean changes, the value of S_n will take on a distribution whose slope is not equal to 0 and the line will drift in the direction of the change. This drift will eventually take the plot outside the control limits signaling a change in the process mean. The calculation of a cumulative sum statistic enables CUSUM charts to distinguish a moderate shift in the mean better

than a Shewhart Chart. This cumulative property, however, also requires that the CUSUM chart be "re-tuned" for the new process mean and restarted each time it signals out of control (Hawkins and Olwell, 1998, p26)

Upper and lower control limits are critical in the responsiveness of the statistical control charts. They are designed to distinguish between usual variation in the process and shifts. They are calculated using a function of the process distribution when the distribution is in control. For Shewhart charts with normal data, the upper and lower control limits are frequently calculated as standard deviations of the batch mean above and below the in control mean. In equation form, the upper/lower control limits are set at:

$$\mu \pm \frac{m\sigma}{\sqrt{n}}$$

(Hawkins and Olwell, 1998, p7)

where m is the number of standard deviations.

As in the example above, a batch of bolts with a mean hardness greater than or less than m standard deviations from the mean will cause an out of control signal on the Shewhart chart. Commonly, control limits are set at 3 standard deviations (m = 3) above and below the correct mean and are referred to as 3 sigma limits. As with the Shewhart charts, CUSUM charts have upper and lower control limits for signaling when the process is out of control. Even though they perform the same function, their calculation and theory is very different. CUSUM control limits are functions of the Average Run Length (ARL) of the chart, the decision interval h, and a reference value k (Hawkins and Olwell, 1998, p32). These three factors, their calculations and their relationships, will be discussed later in section 3.

2. Poisson Univariate Control Charts Methods

Poisson control charts are important because many processes and natural random phenomenon can be better modeled as Poisson rather than Normal, especially when faced with count data (Hawkins and Olwell, 1998, p110, 111). Unless the Poisson rate parameter λ is large, the Shewhart 3-sigma control limits used for normal data are inadequate. This is due to the asymmetry of the Poisson distribution compared to the symmetry of the Normal distribution. For Poisson data, the upper and lower control limits are determined from the probability limits of the Poisson distribution with the given rate λ (Weitzman, 1999, p9).

As stated earlier, CUSUM charts do not plot raw data versus time as do Shewhart charts. For Poisson data when the rate parameter λ is known, CUSUM charts plot cumulative sums of the deviations of the sample values X_i from a reference value k. The upper and lower control limits for each additional data point rely on the previous statistic S_{n-1}, the current data value X_n, and the value of k as shown in the equations:

$$S_n^+ = \max(0, S_{n-1} + X_n - k^+)$$
$$S_n^- = \min(0, S_{n-1} + X_n - k^-)$$

(Hawkins and Olwell, 1998, p112-113)

The values of k^+ and k^- for Poisson CUSUM control charts are functions of the in control mean and the target out of control limits for the mean. The in control mean is the mean of the process being evaluated when the process is considered to be in control. The target out of control limits for the mean are the upper and lower limits for which the process mean is be considered in control. The shifts from the in control mean to the upper and lower limits for the mean are the shifts that CUSUM charts will have the optimal speed of detection.

They are calculated as follows:

$$k^{+} = \frac{\lambda_u - \lambda_o}{\ln(\lambda_u) - \ln(\lambda_o)} \qquad k^{-} = \frac{\lambda_d - \lambda_o}{\ln(\lambda_d) - \ln(\lambda_o)}$$

Where λ_o is the in control mean
λ_d is the out of control mean for a downward shift
λ_u is the out of control mean for an upward shift
(Hawkins and Olwell, 1998, p113).

All previous discussion of control charts has referred to non-self-starting control charts where a large amount of historical data is available. In order for those control charts to be effective, a long period of time is required to collect data when starting a new chart or when "retuning" a CUSUM chart to the new mean after it has detected a shift in the process parameters. This is not attractive to manufacturers who view this "set up" time as a period of no control. Military commanders of units that are the first to deploy to an OOTW environment will not have direct historical data to tune a CUSUM chart. Most unit rotations in Bosnia and elsewhere are typically between six and twelve months. The commanders and their units will most likely rotate out of the environment before they have a time to collect enough data for such charts. CUSUM charts are then only useful to subsequent units if sufficient data has been previously collected and there has not been a change in the process that requires retuning. The volatile nature of OOTW environments, therefore, nearly renders standard non-self-starting CUSUM tools useless to military commanders.

Self-starting control charts enable the user to detect changes soon after implementation of the control charts. They do not require large amounts of historical data to set up and can detect shifts in the process after only a few data points, making them applicable and useful to military commanders in OOTW environments. Weitzman (1999), in his thesis, applied self-starting control chart methodology to a plausibly Poisson process of police use of force. This thesis uses his

methodology for Univariate analysis because, as we shall see, the random nature of enemy incidents in OOTW can be plausibly considered Poisson.

For self-starting Poisson Shewhart charts, the upper and lower control limits are developed by calculating probability limits that are conditioned on the sum of a series of values X_i (Weitzman, 1999, p13). The conditioning argument is based on the property that the Poisson distribution is infinitely divisible and takes the form: $P(X_n = x_n \mid \sum_{i=1}^{n} X_i = S) = \mathit{binomial}\ (S, 1/n)$ (Hawkins and Olwell, 1998, p175).

Weitzman (1999) implemented this formula in Microsoft Excel using the critical binomial value function CRITBINOM(S,p,α). In CRITBINOM, the parameter S is the sum of the preceding n observations, the parameter p is $1/n$ where n is the number of time periods or data batches, and α is the confidence level required. For example, to calculate the upper control limit for the 3rd observation, S would be the sum of these three observations, $p = 1/3$, and α would be a percentage such as .995. This same process is used for the lower control limits except α would be 1 minus the α used for the upper control limit, or 0.005. Using the α's above would produce a 99% confidence interval for the Shewhart control limits of the 3rd observation. It should be noted however, that due to the granularity of discrete functions, an exact 99% confidence interval may not be obtained. The granularity of the discrete functions may produce values close to the target confidence interval, but not exact. For example, discrete function that desires a 99% confidence interval may obtain a 99.2% or a 98.8% confidence interval due to the discrete input values.

The CRITBINOM function, however, requires upper and lower control limit values for the first data point. This thesis uses probability limits, entered by the user, to calculate these initial control limits. The in control test ARL for the first data point depends on the

probability limits, or confidence interval chosen. Although it only affects the ARL of the first point, the choice of probability limits will be discussed in detail, to ensure understanding and maintain consistency throughout this analysis.

The in control ARL, false alarm rate, is derived from the negative binomial distribution when checking for the first error, and which simplifies to a geometric series. In equation form, the in control ARL is solved as follows:

$$ARL_{incontrol} = \frac{1}{1 - prob} \tag{1}$$

where *prob* is the probability limits for the first data point. To obtain a desired in control ARL, this equation can be algebraically manipulated to solve for the appropriate probability limit. For example, if the proper in control ARL is 400, the appropriate probability limit to use is .9975, or 99.75%.

Figure 1 shows an example of a Poisson Self-starting Shewhart control chart using Poisson generated data with a mean of 3. The initial upper and lower control limits were calculated as 7 and 0 using a 99% probability limit. Using the CRITBINOM function to calculate the subsequent control limits allows the limits to change over time as shown. Upward shifts signal a departure if the value is greater than or equal to the upper control limit. Lower shifts, on the other hand, signal a departure if the value is strictly less than the lower control limit. Data point 28 signals a departure because it is plotted on the upper control limit. This enables the user to identify this point as an isolated departure from the mean.

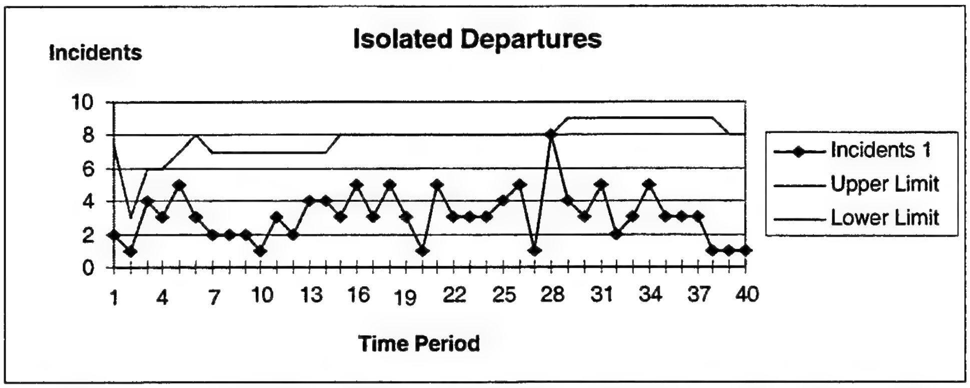

Figure 1. Poisson Self-starting Shewhart Control Chart. Data is generated from a Poisson distribution with a mean of 3. Time periods are measured on the X-axis and number of incidents is measured on the Y-axis. Initial upper and lower control limit values are calculated from Poisson probability limits. Subsequent upper and lower control limit values are calculated using Excel's CRITBINOM function.

For self-starting CUSUM charts where the parameter λ is unknown the CUSUM chart plots the cumulative sum of the deviations of the "transformed" sample values, Y_n, from a reference value k. Using the reference value k, which is calculated as in the non-self-starting CUSUM, and the transformed sample value Y_n, the self-starting CUSUM control limits are calculated as follows:

$$S_n^+ = \max(0, S_{n-1} + Y_n - k^+)$$
$$S_n^- = \min(0, S_{n-1} + Y_n - k^-).$$

This is a slight difference from the non-self-starting CUSUM method but the role of this transformed value, Y_n, is significant and Y_n development demands additional explanation.

For insight into Y_n, assume the process being studied follows a Poisson distribution and the monitored values are discrete count value X_n. Also, assume that the in control mean, λ_o, is unknown. The sample mean, $\overline{X}$, is the appropriate statistic, i.e. maximum likelihood estimator, for estimating λ_o. Now, let W_i = $i\overline{X}$ and condition on W_i which yields X_i~binomial$_i(W_i, 1/i)$. This distribution is parameter free and X_i does not rely on the unknown mean λ_o. Therefore, "if the process

mean shifts from λ_o to λ_1, then the conditional distribution of X_n becomes binomial with a probability $\frac{\lambda_0}{(n-1)\lambda_0+\lambda_1}$" (Hawkins and Olwell, 1998, p175). A change in the process mean will change the probability upward if $\lambda_1 > \lambda_o$ and downward if $\lambda_1 < \lambda_o$. Monitoring the changes in the binomial probability will determine if the mean has shifted up or down.

This conditional distribution for X_n is used to calculate the cumulative probability $A_n = \Pr[Bi(W_n, 1/n) \le X_n]$ (Hawkins and Olwell, 1998, p176). Unlike the continuous case, A_n can only take on a limited number of values because X_n can only assume discrete values 0,1,2,... W_n. The values of A_n are distributed independently even though the values are limited. This can be seen from Basu's lemma (Hawkins and Olwell, 1998, p176).

A_n must now be transformed for use in a CUSUM chart. One point of concern is the cases where $A_n = 1$. This occurs when the initial sequence of X_n's are 0. A_n will equal 1 for the first non-zero X_n. This requires attention in the execution of the transformation.

Transforming X_n to a Poisson variate Y_n with parameter m is done by determining the value of Y_n that minimizes the equation:

$$\left| \sum_{j=0}^{Y_n} \frac{e^{-m} m^j}{j!} - A_n \right|$$ (Hawkins and Olwell, 1998, p177).

In the cases where $A_n = 1$, Y_n is determined by setting $Y_n = X_n$. This transformation is done to get a Y_n that is Poisson with mean m, where m is an estimated process λ. But because of the graininess of the values of A_n brought on by the discrete values of X_n, this is not exactly possible (Hawkins and Olwell, 1998, p177). It is however, very close if the estimated mean is close to the true distribution mean (Weitzman, 1999, p18). The calculation of Y_n in the Poisson self-starting CUSUM

control chart method developed by Hawkins and Olwell is done using a Visual Basic macro developed by them.

Figure 2 shows a Poisson self-starting CUSUM control chart using the same generated Poisson data as in Figure 1 with mean equal to three. The upper and lower control limits were calculated using Fortran based software package ANYGETH.exe with an average run length (ARL) of 100. ANYGETH.exe and ARL's will be discussed in detail in the next section, section 3.

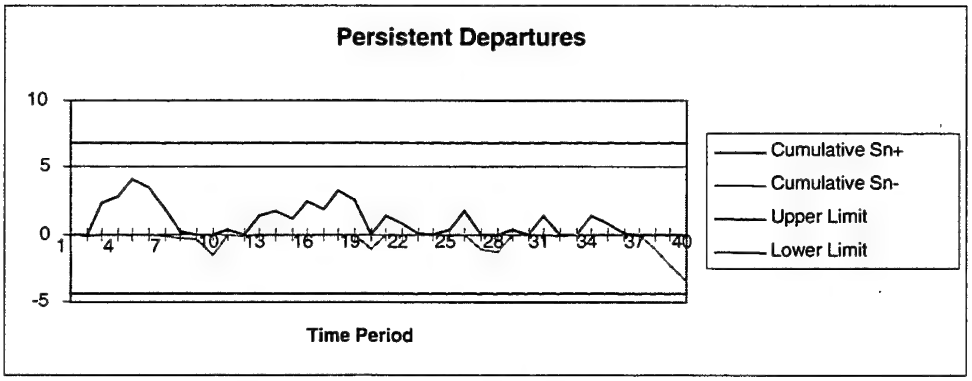

Figure 2. Poisson Self-starting CUSUM Control Chart. Data is generated from a Poisson distribution with a mean of 3. Time periods are measured on the X-axis and the calculated values of the cumulative statistics S_n^+ or S_n^- are measured on the Y-axis. The target in control mean is 2.95. The out of control mean for an upward shift is 4.4 and the out of control mean for an downward shift is 1.5. The control limits are set at 6.8 for an upward shift and -4.4 for a downward shift. The average run length (ARL) is 100.

3. Average Run Length and CUSUM Control Chart Limits

Poisson self-starting CUSUM charts require five parameters before they can be run. The five parameters are the average run length (ARL), the upper and lower control limits (h^+ and h^-), and the reference values (k^+ and k^-) (Hawkins and Olwell, 1998, p44). These parameters are interrelated and can be calculated using available computer packages such as ANYGETH.exe and ANYGETARL.exe. Using a software package such as ANYARL.exe allows one to calculate the associated ARL with a given k and h, where the software package ANYGETH.exe calculates the upper and lower control limits given a k and an ARL. It is common to select the ARL

based on the discussion below and calculate the reference value k from the target in control and out of control means. ANYGETH is then used to solve for the upper and lower control limits. Directions for the software package ANYGETH developed by Hawkins and Olwell, which is used in this thesis, are listed in Appendix C.

The ARL for a chart is defined as the expected number of time periods (runs) before the chart signals a shift when in fact none has occurred (Montgomery, 1985, p287). It is commonly referred to as the average time between false alarms. It is important to note that there is a trade off when determining the ARL that is analogous to the trade off between Type I and Type II error in classical hypothesis testing. In hypothesis testing, reducing the amount of Type I error increases the amount of Type II error in the test. In CUSUM charting, increasing the ARL decreases the number of false alarms that the chart will signal, but it also increases the time required by the CUSUM to detect a shift. Decreasing the ARL increases the number of false alarms, but decreases the time required to detect a shift (Hawkins and Olwell, 1998, p33). The choice of the proper ARL depends on the concerns of the decision-maker and the costs associated with a false alarm and a missed shift in the process.

Many manufacturing processes use ARL's higher than 1000 because the costs associated with a false alarm, which often include shutting down the process, can be enormous compared to harm of producing a improper product. Take for example a production line of the Ford Motor Company that produces 10 sport utility vehicles an hour. Ford receives a profit of $10,000 per vehicle. Managers may use a high ARL when checking the vehicles for defective window seals. The cost associated with not detecting a defective window seal, repair at the dealership, is small compared to the cost of shutting down the assembly line for an hour because of a false alarm, $100,000. On the other hand, managers may use a small ARL when checking for defective brakes. In this case,

the costs of shutting down the assembly line for an hour, $100,000, is small compared to the recall of vehicles and potential liability costs (economic and human) associated with an accident caused by a faulty brake mechanism.

It is important to note that ARL's used in combined tests have an additive affect on the overall process ARL. Combined tests are any tests used simultaneously on a data set. Upper and lower control limits are an example of two tests that when used together constitute combined tests. For example, if ARL's of 100 are used in 2 combined tests, say an upper and a lower control limit, then the combined test can be expected to produce 2 false alarms, 1 for each limit, in 100 periods. The process ARL is therefore 2 in 100, or 1 in 50, not 1 in 100. Van Dobben de Bruyn (1968) showed that for combined systems, a conservative method of calculating the test ARL's to achieve the proper overall ARL is as follows:

$$\frac{1}{ARL_{combined}} = \sum \frac{1}{ARL_{test}} \qquad (2)$$

(Hawkins and Olwell, 1998, p55). This thesis uses different test ARL's in order to achieve an overall or combined ARL of 100 for each type of analysis. The individual univariate analysis of the three separate data categories has four tests: Shewhart upper control limit, Shewhart lower control limit, CUSUM upper control limit, and CUSUM lower control limit. A test ARL of 400 is used for each of these four tests in order to obtain a combined ARL of 100 for each individual data category.

Multivariate analysis uses a total of 16 tests. From equation 2, a test ARL of 1600 is desired to obtain a combined ARL of 100. 12 of the 16 tests in the multivariate analysis use an ARL of 1600. However, four tests in the nonparametric multivariate analysis use confidence intervals for the upper and lower control limits. These confidence intervals affect the in control ARL's similar to the probability limits explained above. Using an ARL of 1600 in Equation 1 and solving for the

confidence interval, results in a confidence interval of 99.9375%. Rounding this confidence interval to 99.94% for simplicity altered the ARL to 1667. This is however a sufficiently close approximation to the desired ARL of 1600. A detailed discussion of the different ARL's used in multivariate analysis and their calculations are explained in Chapter V, Section A2. The combination of these different ARL's using equation 2 resulted in an overall combined ARL of 101.015 for the multivariate analysis, which is sufficiently close to 100.

The methods used in calculating the ARL's or upper and lower control limits in CUSUM charts, including those used in computer packages, take three common forms: solving integral equations, solving discrete Markov chain approximations to the integral solution, and using simulation (Hawkins and Olwell, 1998, p153).

The integral equation for continuous variables is as follows: $L(z) = 1 + L(0)F(k-z) + \int_0^h L(x)f(x+k-z)dx$ for each $z \in (0,h)$ (Hawkins and Olwell, 1998, p154). *L(z)* is the average run length for the CUSUM that starts at $S_o = z$. The first component of this equation is the probability that the chart will test another value. This value is 1 because at least one more observation is always drawn for $z \in (0,h)$. The second component, *L(0)F(k-z)*, is the probability that the next data value returns the CUSUM to zero *(F(k-z))*, multiplied by the average run length from zero *(L(0))*. The final component "is the integral of the average run length for the next value of the CUSUM if it is between 0 and *h*, multiplied by the probability that this next value occurs" (Hawkins and Olwell, 1998, p154).

The software package ANYGETH uses the discrete Markov chain approximation to the integral solution to solve for the upper and lower control limits. The discrete Markov chain approximation to the integral solution solves the discrete analog of the integral equation above.

This analog takes the form $L(z) = 1 + \sum_{i=0}^{M} L(i)R_{i,z}$, where $R_{i,z}$ is the Markov transition matrix not including transitions to and from the last state. The last state is not included because the ARL from State M+1 is always zero (Hawkins and Olwell, 1998, p155). The Markov equation in matrix form is as follows: $(I - T)\lambda = 1$, where I is an identity matrix, T is the transition probability matrix, λ is a vector of length M+1 of ARL values for CUSUM's starting in the corresponding state, and 1 is a M+1 vector whose values are all 1. Solving the equation results in the appropriate ARL for the given h and k (Hawkins and Olwell, 1998, p155). Because they are interrelated, ANYGETH solves for the value of h given an ARL and k.

The third method, simulation, involves simulating the process used to calculate the CUSUM, determining and recording the run lengths, and averaging the run lengths to determine the ARL. Although work has been done in improving the precision of the estimates for the ARL's, simulation remains an intensive and inefficient method (Hawkins and Olwell, 1998, p156). In this thesis, simulation is not used to calculate the ARL. Instead simulations are used to verify the theory and software developed in this thesis. Simulations, run multiple times using generated data sets with known parameters, verify the accuracy of the resulting CUSUM charts.

4. Discussion of CUSUM Optimality

CUSUM methods have been shown to possess various optimality properties. In the context of Statistical Process Control, optimality is reserved for the scheme that is quickest to detect a shift in the process from in control to out of control. "Or more formally, among all procedures with the same in-control ARL, the optimal procedure has the smallest expected time until it signals a change, once the process shifts to the out-of-control state" (Hawkins and Olwell, 1998, p138).

Moustakides (1986) proved that CUSUM charts are optimal in this sense. "Among all tests with the same in control ARL, CUSUM has the smallest expected run length out of control" (Hawkins and Olwell, 1998, p138). CUSUM charts are however "tuned" for a specific shift in a specific distribution, and therefore, the CUSUM is optimal for detecting only this specific shift. A different CUSUM would be optimal for detecting other shifts. This would greatly diminish the applicability of CUSUM charting, if it were not for the robust performance of CUSUM. CUSUM charts are robust in that the optimality qualities nearly hold for shifts close to that which it was designed to detect. "That is to say, while the CUSUM for detecting a one-standard-deviation shift is only optimal diagnostic for that particular shift, it does nearly as well as the optimal CUSUM for all shifts "not too far" from one standard deviation" (Hawkins and Olwell, 1998, p139).

The robustness of CUSUM charting methodology can be checked by comparing the out of control ARL's calculated by ANYGETH.exe for a targeted shift to those calculated by ANYGETH.exe for a nearly equivalent shift using the same ARL and the same reference value k. For example, a process with a target in control $\lambda_0 = 3$ and an out of control $\lambda_u = 6$ will result in ANYGETH.exe returning an exact reference value of $k = 4.328$. In this example, the exact reference value of $k = 4.328$ is rounded to a value of $k = 4.4$. Using an ARL of 100, ANYGETH.exe calculates an in control ARL of 116.07 and an out of control ARL of 3.5. Running ANYGETH.exe again with the same in control $\lambda_0 = 3$, the same rounded value of $k = 4.3$, and the same ARL of 100, but with an out of control $\lambda_u = 5$, the resulting in control ARL = 116.07 and the resulting out of control ARL = 6.

Because both executions of ANYGETH.exe use the same in control $\lambda_0 = 3$, the same rounded value of $k = 4.3$, and the same ARL, they are both tuned to optimally detect a shift from $\lambda_0 = 3$ to $\lambda_u = 6$. The in control ARL's

are the same because tuning the charts for the same shift results in the same false alarm rate. However, the out of control ARL's are slightly different because the out of control ARL's are the measure of how quickly the CUSUM charts detect the shift in the process (Hawkins and Olwell, 1998, p36). The out of control ARL for a shift of $\lambda_u = 5$ is larger than the out of control ARL for the shift of $\lambda_u = 6$ meaning that it will take longer for the charts to detect the smaller shift than the larger shift. The robustness of the CUSUM charts is evident here in that even though the charts were not specifically tuned for the shift of $\lambda_u = 5$, they will none the less detect the smaller shift. The charts require additional time to detect the smaller shift. This detection time difference is the difference between the two out of control ARL's, or 2.5 time periods. Depending on the situation, this difference is minimal. Users can therefore capitalize on the robustness of CUSUM charting and apply them with confidence knowing that the charts, although not optimal, are nearly so.

B. MULTIVARIATE CONTROL CHART METHODS

Multivariate control charts are used to analyze a collection of process measurements, not just one measurement as in the univariate control chart methods described earlier. Two major benefits of multivariate control charts are that they are more sensitive to multiple shifts than are univariate control charts used individually and they also improve the diagnostics of the shifts. Better diagnosis of the nature of the change will enable managers to better identify and fix the cause of the shift. Using a published example, the quality of coal produced from a washing plant is judged based on the yield and the ash content of the coal after it has undergone the washing process. Two factors that influence the final product are the effectiveness of the washing process and the quality of raw coal that was used in the process. If a shift occurs in the amount of ash in the produced coal,

univariate control charts will detect the shift and may attribute the shift to a change in the washing process. It may in fact be a result of a change in the quality of the raw coal shipment used. Multivariate control charts will detect the shift and help attribute the cause of the shift to the correct cause. In the above example, multivariate control charts would attribute the shift to the quality of coal used and prevent the managers from searching for a problem in the process (Hawkins and Olwell, 1998, p190).

The Normal distribution is the basis for much statistical work done with multivariate data. This is a result of the Normal distribution having preferred statistical properties and because, for multivariate work, there are "few other manageable widely know distributions available" (Hawkins and Olwell, 1998, p191). One of the more favorable properties of the multivariate normal distribution is that its marginal distributions and conditional distributions are also normal. It is also useful to know that linear combinations of multivariate normal variates are also normally distributed (Anderson, 1984, p24). In general, the multivariate normal distribution has often been found to be a sufficiently close approximation to the analyzed population, justifying its use (Anderson, 1984, p4). These favorable properties, as well as others, do not usually hold for other distributions, making multivariate normal the distribution of choice.

We will use the following parameterization in our mulitvariate analysis. p is the number of related measurements taken and X_n is the n^{th} sample of the p-component process measurement. The multivariate normal assumption then states that the vectors X_n will follow a common multivariate normal distribution with a mean vector μ and a covariance matrix Σ. In equation form: $X_n \sim N(\mu, \Sigma)$ (Hawkins and Olwell, 1998, p191). The covariance matrix Σ is the key factor in capturing the relationships between the different process measurements made on the same sample and

is responsible for benefits of multivariate control charts over univariate control charts. If the process measurements are uncorrelated, the off diagonal elements of the covariance matrix will be zero. In this case, it may seem that multivariate control charts are no better than a collection of univariate control charts. This however is not entirely true, in that multivariate control charts may still offer better insight if the cause of a shift effects the multiple properties measured (Hawkins and Olwell, 1998, p191). It is important to note that the model assumes that the in control X_n vectors are independent for different n. That is to say that although the p-measurements taken from sample n may be correlated, they are independent from the p-measurements taken in sample n+1. It is also important to note that the measurements in the X_n vector must relate to the same product, not necessarily the same time (Hawkins and Olwell, 1998, p191-192). In the coal washing example, if two measurements are being taken on a given sample of coal, one before it is washed and one after it is washed, the observer must ensure that the before washing measurement stays linked with the after washing measurement of the same batch of coal. If the measurements were taken at the same time, then the before washing measurement and the after washing measurement would come from different batches of coal and would be meaningless.

In graphical terms it is clear to see the actions of the multivariate methods. Using the coal washing example, if the yield of the washed coal is plotted against the ash content of the washed coal, the plot will assume some form of a bivariate distribution depending on the correlation between the two variables, as shown below:

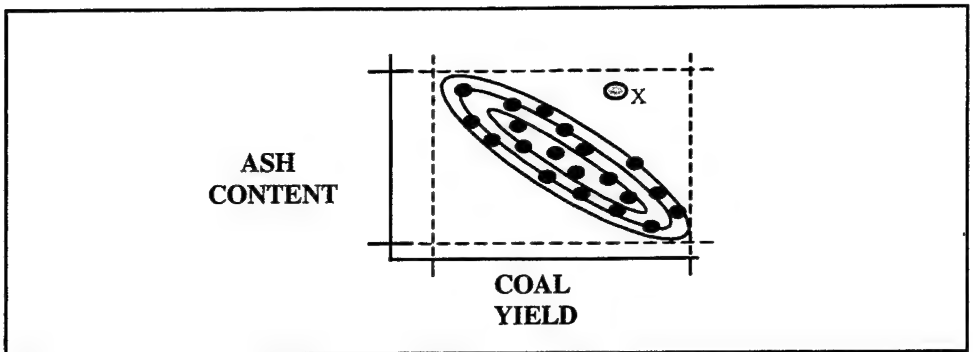

Figure 3. Graphical Depiction of Multivariate Methods. Measurements of coal yield per shipment on the X-axis against the corresponding ash content of the shipment on the Y-axis. The data point X lies in the range of both ash content and coal yield, but is an outlier to the bivariate distribution of the data.

From Figure 3, it is clear that the data point "*X*" does not follow the bivariate distribution of the other samples. This difference of sample "*X*" from the other samples may be caused by an increase in coal quality that offsets a decrease in the effectiveness of washing process on that sample. Multivariate methods will detect this difference and will signal a shift in the process from in control to out of control. The data point "*X*" may not signal a shift in Univariate methods. It lies inside the range of ash content and coal quality, and therefore may be inside the separate control limits for each variable.

For multivariate normal Shewhart control charts, Hotelling's T^2 statistic is the most powerful test statistic. This assumes that the *p*-component vector X is multivariate normal, $X_n \sim N(\mu, \Sigma)$, and that Σ is known. The preferred Hypothesis test is to test the null hypothesis H_o: $\mu = \mu_o$, against the alternate hypothesis H_a: $\mu \neq \mu_o$. This test is targeted at any shift in μ, and from multivariate theory, the most powerful affine invariant test statistic for H_o against H_o rejects the null hypothesis if the value of T^2 is large. T^2 is calculated as follows:

$$T^2 = (X_n - \mu_{o_{N-1}})' \sum\nolimits_{n-1}^{-1} (X_n - \mu_{o_{n-1}})$$

and is compared to the Chi Squared with p degrees of freedom, or $T^2 \sim \chi_p^2$ (Hawkins and Olwell, 1998, p192).

Affine invariant tests are test statistics that are "unaffected by a full rank linear transformation of the vector X", i.e. $Y = AX$ (Hawkins and Olwell, 1998, p192). The restriction to affine invariant tests is used when the possible shift of μ is unknown. If there is knowledge about the type of shift in μ that might occur, the affine invariant restriction can be discarded. The hypothesis test now used will test the null hypothesis H_o: $\mu = \mu_o$, against the alternate hypothesis H_a: $\mu = \mu_1$. This test statistic for H_o against H_a is $Z = (X - \mu_o)'\sum^{-1}(\mu_1 - \mu_o)$. Z follows the normal distribution shown below with $\lambda = m\Delta'\sum^{-1}\Delta$ where Δ is the size of the shift in the mean:

$$\begin{array}{ll} Z \sim N(0,\lambda) & \mu = \mu_o \\ Z \sim N(\lambda,\lambda) & \mu = \mu_1 \end{array}$$

(Hawkins and Olwell, 1998, p192).

This is a significant improvement over the T^2 test because it essentially shows the test where to look for a shift. Also, the improvement this test makes over the T^2 test gets greater as p gets larger (Hawkins and Olwell, 1998, p193-194). This method is presented to increase understanding of the material. This thesis did not consider this method in analyzing the SFOR data set because there is no information or knowledge about the type of shift that might occur.

In multivariate CUSUM control charts, as in univariate CUSUM control charts, the issue of detecting smaller but persistent shifts in the data still requires a method that accumulates information across successive observations. The univariate recursion to address this issue is as listed earlier:

$$S_n^+ = \max(0, S_{n-1} + X_n - k^+)$$
$$S_n^- = \min(0, S_{n-1} + X_n - k^-).$$

In the multivariate case, however, a vector $\boldsymbol{X_n}$ replaces the scalar X_n. The best application of this vector in the Univariate recursion is unclear (Hawkins and Olwell, 1998, p195).

Crosier (1988) introduced a multivariate CUSUM method that accumulates on the scale of the vector $\boldsymbol{X}$. Accumulating on the vector $\boldsymbol{X}$ initializes the CUSUM vector $\boldsymbol{S_n}$ to a zero vector and alleviates the problem of when the shift is in a direction other than that proposed. The appropriate recursion is as follows:

$$S_n = \begin{cases} 0 & \textit{for } C_n \le k \\ \dfrac{S_{n-1} + X_n - \mu_o}{1 - k/C_n} & \textit{for } C_n > k \end{cases}$$

where $C_n = (S_{n-1} + X_n - \mu_o)' \sum^{-1} (S_{n-1} + X_n - \mu_o)$ (Hawkins and Olwell, 1998, p195).

Note: C_n, S_n, S_{n-1}, μ_o are vectors, Σ^{-1} is a matrix

This recursion causes the CUSUM to signal if $S_n' \Sigma^{-1} S_n$ is greater than the scalar decision interval h. This recursion uses the T^2 metric for its final decision. "It has no known optimality properties, but does appear to have good practical purpose" (Hawkins and Olwell, 1998, p196).

C. DEVELOPED THEORY OF THE NONPARAMETRIC MULTIVARIATE CONTROL CHART METHODS

1. Theory

As stated above, the multivariate Normal distribution forms the basis for typical multivariate control chart methods. The multivariate normal distribution has robustness for other distributions, but the robustness depends on assumptions between the multivariate normal and the specific distribution of the process. This thesis chose to initially model the SFOR Incident Data as Poisson. The Poisson distribution was chosen because the incidents of enemy actions in OOTW are uncoordinated and stochastic counts, making them plausibly Poisson. Multiple tests, shown in Appendix D, verified that the data could be considered Poisson. But because there is not a commonly accepted model

for multivariate Poisson data, nor is there a multivariate scheme for Poisson data, this thesis chose to use nonparametric techniques for the multivariate control chart analysis. A nonparametric method will forego any need for assumptions about the data being Poisson or any need for multivariate Normal approximations to the multivariate Poisson. In effect, nonparametric techniques will be applicable to all data sets regardless of the underlying distribution (Anderson, 1984, p5).

The multivariate analysis method developed in this thesis consists of two parts. First, univariate analysis is conducted simultaneously on the three data categories and will be referred to as simultaneous univariate analysis to avoid confusion between it and the individual univariate analysis. Second, a nonparametric permutation technique, developed in this thesis and described in detail below, is conducted to analyze the multivariate aspects of the data categories. This will be referred to as nonparametric multivariate analysis. The crucial concept in these two parts of the multivariate analysis method is that a persistent departure in any one of the CUSUM charts, simultaneous univariate CUSUM charts or the nonparametric multivariate CUSUM charts, requires that all charts be retuned and restarted at the originating time of the detected shift. This is done to maintain the time relationship of the data categories and to maintain the correlation between the data categories.

Simultaneous univariate analysis is similar to individual univariate analysis as previously explained except for two key issues. As stated above, the simultaneous univariate analysis control charts, as well as the nonparametric multivariate control charts, must be retuned and restarted when a persistent shift is detected in any of simultaneous univariate CUSUM control charts or the nonparametric multivariate CUSUM control chart. Also, the combined ARL in the analysis is now dependent on the 16 different tests contained in the simultaneous univariate analysis and the nonparamteric multivariate analysis. The 16 tests are

as follows: upper and lower control limits for each data category in the simultaneous univariate Shewhart control charts, upper and lower control limits for each data category in the simultaneous univariate CUSUM control charts, an upper and a lower control limit in the nonparametric multivariate Shewhart control chart, and an upper and a lower control limit in the nonparametric multivariate CUSUM control chart. Calculating the appropriate ARL's for these 16 tests in order to obtain the correct combined ARL is explained in detail in Chapter V, section A2, Multivariate Parameters.

The nonparametric permutation technique developed for the nonparametric multivariate analysis of the data extends common distribution free based methods and applies it to multivariate control charts. This technique begins by taking numerous permutations of the data. For each permutation, the T^2, S_n^+, and S_n^- statistics, from equations 3, 4, and 5 below, were calculated for each time period and then stored in separate arrays for each time period. After all permutations have been conducted, each array is sorted from lowest to highest. The upper and lower control limits for each time period is calculated from this ordered array of permutated statistics. For example, after taking 1000 permutations of the data, each time period will have three corresponding arrays of 1000 T^2 statistics, S_n^+ statistics, and S_n^- statistics. The arrays are sorted from lowest to highest and for a 99% confidence interval, the 0.5% and 99.5% percentile values in the arrays are used as the upper and lower control limits for each time period. The control limits for the multivariate Shewhart charts use the T^2 statistic. The upper control limit for the multivariate CUSUM charts use the S_n^+ statistic where as the lower control limit for the multivariate CUSUM charts use the S_n^- statistic.

As stated above, multivariate Shewhart control charts' upper and lower control limits are established from the distribution of the T^2

statistic for a two-sample problem. This T^2 statistic tests the null hypothesis that the mean of the first normal population is equal to the mean of the second population and the covariance matrices are equal but unknown. In this test, T^2 is calculated as follows:

$$T^2 = \frac{N_1 N_2}{N_1 + N_2}(X_n - \overline{X}_{n-1})' \sum\nolimits_{n-1}^{-1}(X_n - \overline{X}_{n-1}) \tag{3}$$

where: N_1 is the number of samples in the 1^{st} population
N_2 is the number of samples in the 2^{nd} population
X_n is the observation at time period n
$\overline{X}_{n-1}$ is the average of the observations up to time period n-1
$\Sigma^{-1}{}_{n-1}$ is the inverse covariance matrix at time period n-1.

Under the assumption of normality, it is distributed as T^2 with $N_1 + N_2 - 2$ degrees of freedom and the critical region is:

$$\mathrm{T}^2 > \frac{(N_1 + N_2 - 2)p}{(N_1 + N_2 - p - 1)} F_{p, N_1 + N_2 - p - 1}(\alpha)$$ (Anderson, 1984, p167).

In order to make this a self-starting test, this thesis calculated the T^2 Statistic iteratively, testing if the next observation in the sample data is statistically similar to the mean and covariance of the previously observations. For example, on the 5^{th} permutation, the covariance matrix of the data and the means of the variates are calculated for the first four observations. N_1 is equal to four, N_2 is always equal to one, X_n is the fifth sample observation, $\overline{X}_{n-1}$ is the mean of the first four observations, and $\Sigma^{-1}{}_{n-1}$ is the inverse covariance matrix of the first four observations. Such a step is done for each data observation after an initial start up time. The initial start up time is required to be at least as many periods as the number of data variates you are analyzing in order to obtain a non-singular covariance matrix. Using three data variables, simulations revealed that start up periods of 4, 5, and 6 resulted in near singular covariance matrices and extreme values of T^2 which skewed the graphs considerably. Using 7 periods for the start up time was sufficient to avoid this issue.

The chart in Figure 4 is a plot of the calculated T^2 statistic from generated multivariate Poisson data versus the appropriate F values, based on an assumption of normality. The graph shows numerous upward and downward transient shifts, or departures, in the process when in fact there should be none. The misleading nature of this graph clearly shows that assuming normality is not the correct method to use.

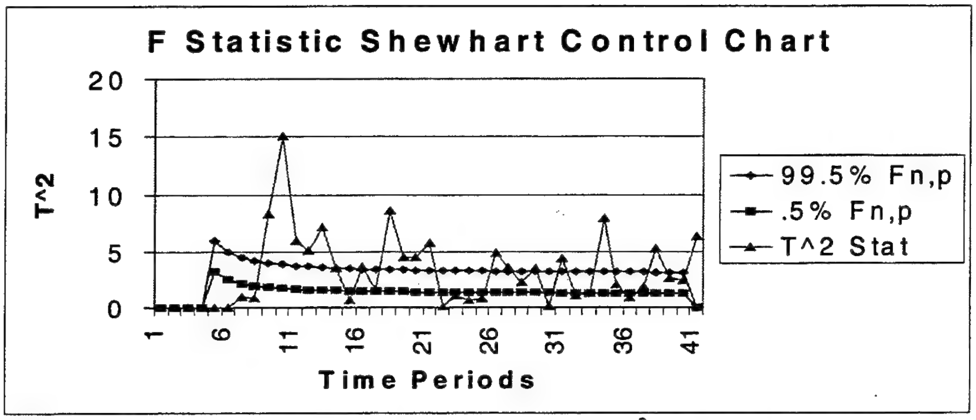

Figure 4. Shewhart Control Chart of T^2 vs F Distribution. Multivariate Poisson generated data with mean equal to 3. Time periods are measured on the X-axis and the values of the calculated T^2 statistics are measured on the Y-axis. Upper and lower control limits are derived using the 99.5% and .5% values of the F distribution.

In an attempt to improve this control chart, the nonparametric permutation technique discussed above was used to get the 99% confidence interval of the T^2 statistic from equation 2 for each sample period. When these were used as the upper and lower control limits, the graph better reflected the consistency of the data with no isolated departures as shown in Figure 5.

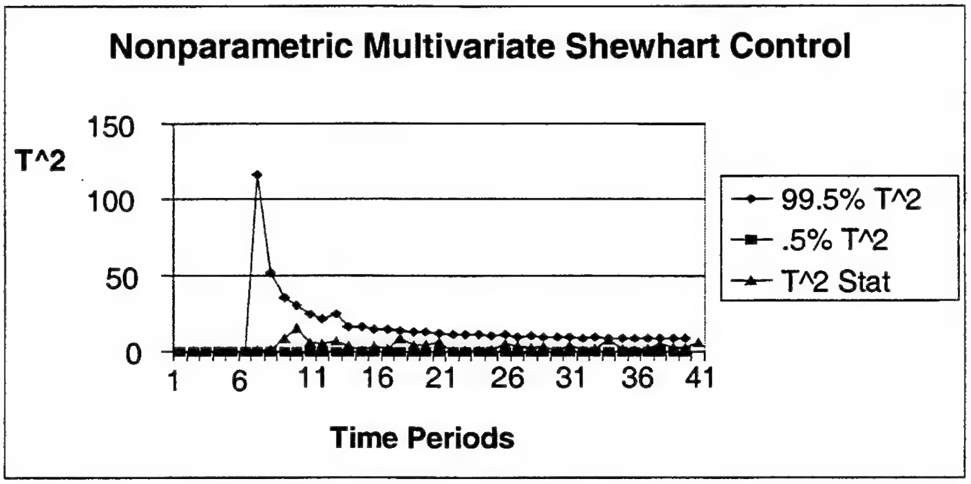

Figure 5. Nonparametric Multivariate Shewhart Control Chart Without Departure. Multivariate Poisson generated data with mean equal to 3. Time periods are measured on the X-axis and the values of the calculated T^2 statistics are measured on the Y-axis. Upper and lower control limits are derived using the nonparametric permutation technique.

Applying the nonparametric permutation technique with a 99% confidence interval to a data set containing an isolated departure at time period 37 is shown in Figure 6. The chart signals an isolated upward departure at time 37.

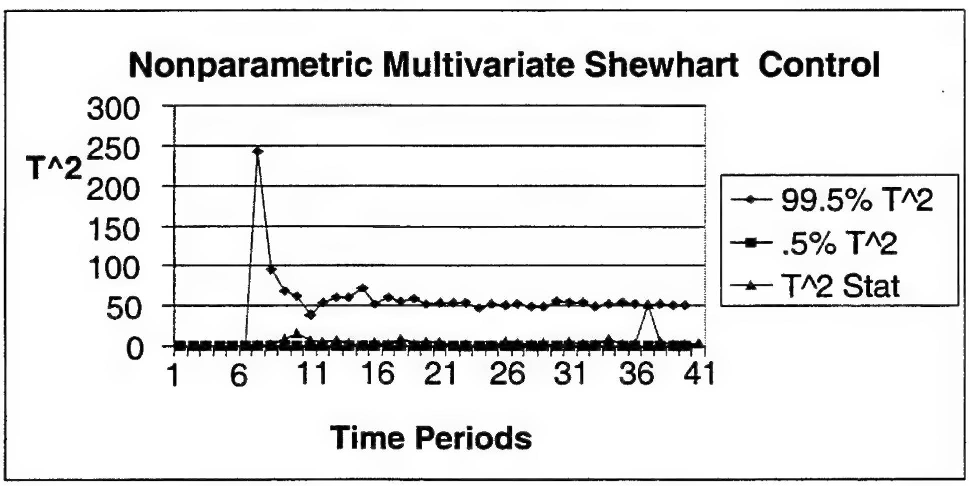

Figure 6. Nonparametric Multivariate Shewhart Control Chart With Departure. Multivariate Poisson generated data with mean equal to 3. Time periods are measured on the X-axis and the values of the calculated T^2 statistics are measured on the Y-axis. Upper and lower control limits are derived using the nonparametric permutation technique. An isolated upward departure is detected at time period 37.

This graph signals the upward departure at time 37 as expected. The chart plots subsequent time period observations inside the control limits verifying that this is an isolated departure in the data.

We created the isolated departure by viewing the data in a 3-dimensional graph and then inserting a point that lies outside the data's multivariate contours. The 3 dimensional graph of the data set with the outlier inserted is shown below in Figure 7.

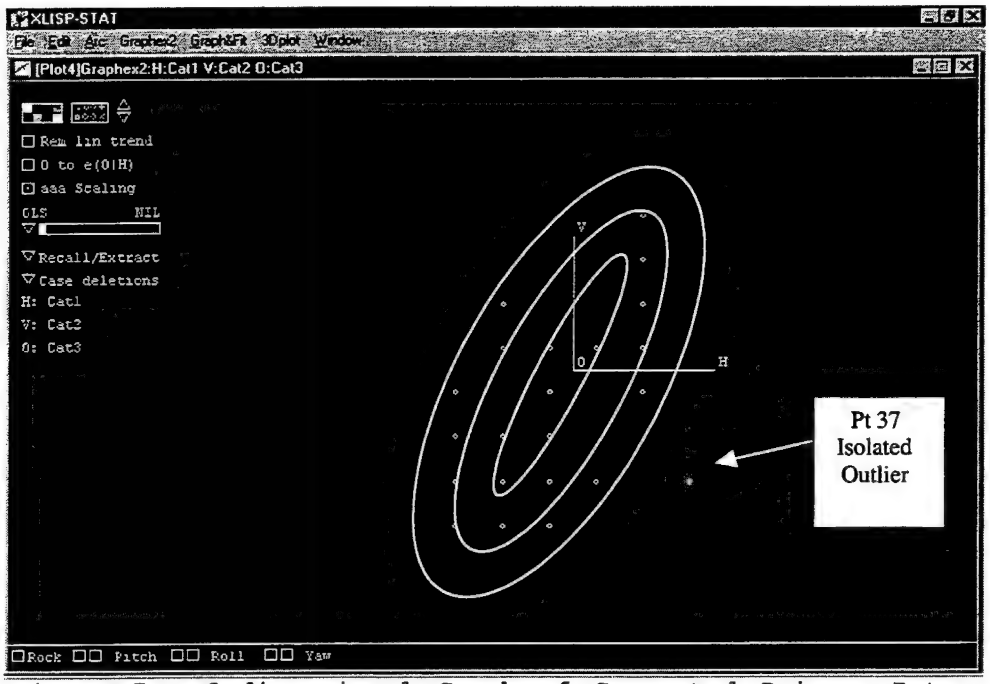

Figure 7. 3-dimensional Graph of Generated Poisson Data. The mean of the Poisson data is 3. To create the isolated departure, a multivariate data point that lies outside the data's multivariate contours was inserted at period 37.

For the self-starting nonparametric multivariate CUSUM, the upper and lower control limits were calculated from a 99% confidence interval of the permutated S_n^+ and S_n^- as shown:

$$S_n^+ = \max(0, S_{n-1} + T_n^{\,2} - k^+)$$
$$S_n^- = \min(0, S_{n-1} + T_n^{\,2} - k^-) \cdot$$

There is no current theory for the calculation of multivariate nonparametric reference values. It can be shown from the equations, however, that the reference values, (k^+ and k^-), affect the slope of the

upper and lower control limits and should be close to the corresponding average values of T^2.

If they are not close to the average value of T^2, the upper and lower control limits will converge either on zero, $+\infty$, or $-\infty$. As seen in Figure 8, for example, if the reference value k^+ is too large, the upper control limit will converge towards zero because, on average, you will continually subtract much more than the current value of T^2.

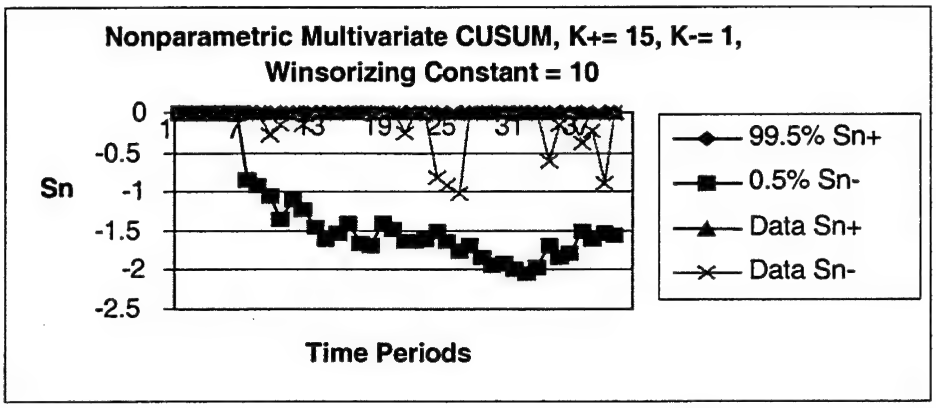

Figure 8. Nonparametric Multivariate CUSUM Control Chart Where k^+ is too Large. Time periods are measured on the X-axis and the calculated values of the cumulative S_n^+ and S_n^- statistics are measured on the Y-axis. The upper and lower control limits are calculated using the nonparametric permutation technique. Large k^+ causes upper control limit to converge on zero.

If the reference value k^+ is too small, as shown in Figure 9, the corresponding control limit will diverge away from zero because, on average, you will continually add more than the current value of T^2.

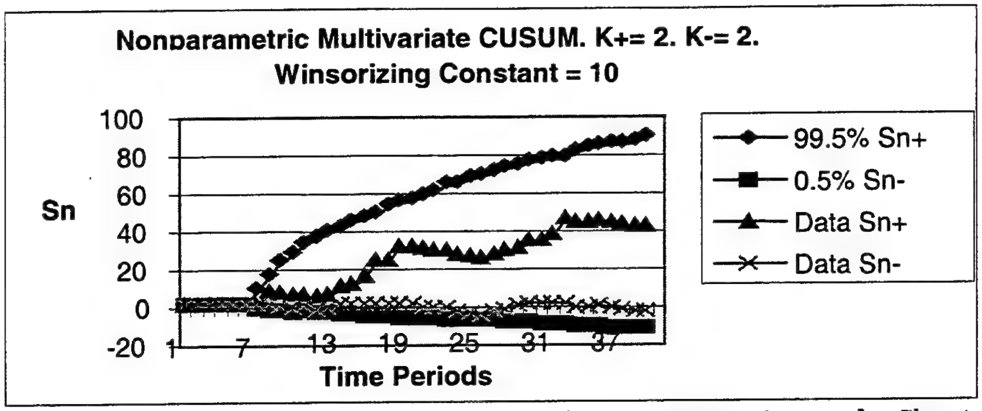

Figure 9. Nonparametric Multivariate CUSUM Control Chart Where k^{+} is too Small. Time periods are measured on the X-axis and the calculated values of the cumulative S_n^+ and S_n^- statistics are measured on the Y-axis. The upper and lower control limits are calculated using the nonparametric permutation technique. Small k^+ causes upper control limit to diverge from zero.

Similar but opposite effects occur with the reference value k^-. If the value of k^- is too large, the lower control limit will converge on $-\infty$ and if k^- is too small the lower control limit will converge on zero. This thesis used multiple simulations to fine tune the reference values until one was found that produced suitable control limits.

Once these control limits are determined, the values of S_n^+ and S_n^- calculated from the original data observations were plotted against these upper and lower control limits. The results are shown in Figure 10. In this case, the process is constant with mean equal to three, k^+=3.75, k^-=1, and a Winsorizing constant (explained below) equal to 10. The reference values k^+=3.75 and k^-=1 produced upper and lower control limits that stabilize near 30 and -1. The nonparametric permutation technique correctly shows a process in control with no signaled shifts in the process.

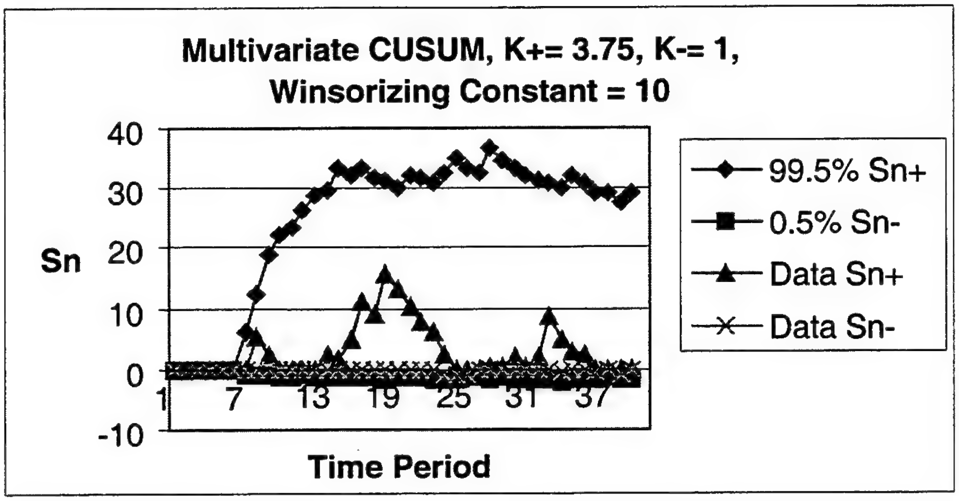

Figure 10. Nonparametric Multivariate CUSUM Control Chart Without Shift. Multivariate Poisson generated data with mean equal to 3. Time periods are measured on the X-axis and the calculated values of the cumulative S_n^+ and S_n^- statistics are measured on the Y-axis. The upper and lower control limits are calculated using the nonparametric permutation technique. Suitable values of k^+ and k^- causes upper and lower control limits to converge on a nonzero value. The process is in control.

When a shift in the covariance structure is added to the process, a shift is signaled in the chart as shown in Figure 11. The shift signals at time period 39. Upon close analysis of the graph, the shift appears to start at time period 38, which is the first "shifted point" after the last time period that the "Data S_n^-" line leaves the X axis before exceeding the control limit. Time period 38 was in fact when the change to the covariance structure was added.

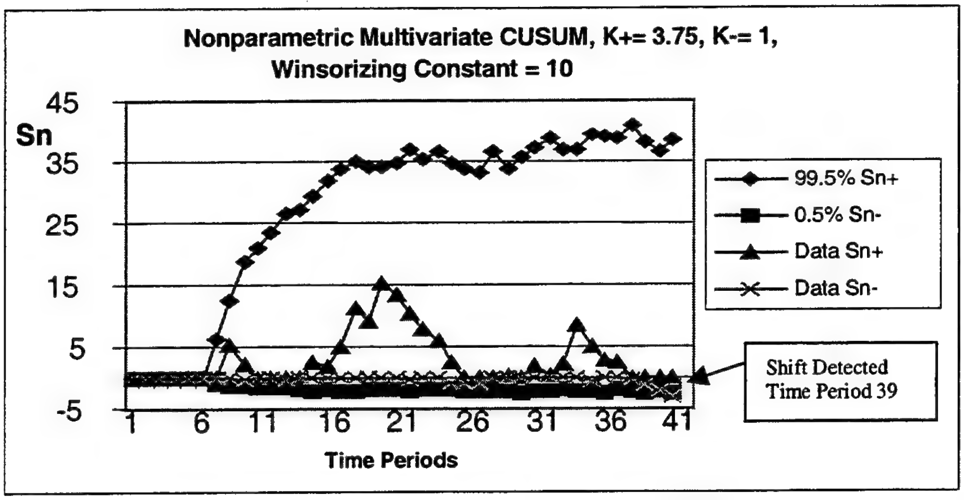

Figure 11. Nonparametric Multivariate CUSUM Control Chart With Shift. Generated multivariate Poisson data with mean equal to 3 and a shift in the covariance structure of the data at time period 38. Graph signals a downward shift at time period 39.

The change to the data set that caused this downward shift in the graph is a change in the variability of the data towards the mean. In other words, the covariance of the data is decreasing. Having all the data observations after time period 37 equal the mean of 3 produced this shift. Graphically this shift can be depicted as in figure 12.

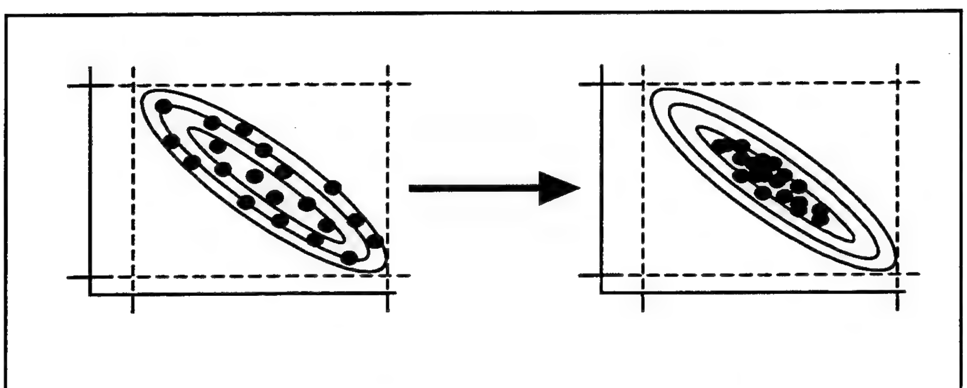

Figure 12. Graphical Depiction of a Decrease in the Covariance Structure. Plotted point fall closer to the center contour line of the bivariate distribution.

This reduction in the covariance structure will signal a departure in multivariate CUSUM charts as shown in Figure 11, but will not cause a shift in the univariate charts. This demonstrates a strength of multivariate analysis.

The downward shift in Figure 11 is difficult to see because of the near zero values of the S_n^- statistic and the lower control limit. In Excel, the graphs can be expanded to simplify the identification of a departure and the time period in which it started. To further simplify the identification of a departure, the Excel program "Multivariate" developed in this thesis identifies a shift as "hot" in text boxes corresponding to the time period of the detection on the Excel worksheet "data1". An example of the Multivariate Excel worksheet "data1" and the text boxes denoting a shift is shown in Figure 13.

An initial start up period is also required for the CUSUM charts, but the start up period must be longer than in Shewhart charts. Additional periods are required for the CUSUM charts in order to avoid "near" singular covariance matrices in the calculation of the T^2 statistic. Such near nonsingular covariance matrices early in the permutation process will produce extreme values of T^2. Because the CUSUM charts are cumulative by nature, these initial extreme values T^2 will skew the remaining values of T^2 resulting in an incoherent graph. By setting the required start period for the trivariate examples used for the graphs above at 7, this problem was avoided.

Another point of concern based in the cumulative nature of the CUSUM chart is the effect a single large T^2 statistic has on the CUSUM chart. A single large value of the T^2 statistic is considered an isolated value of T^2. This should cause a signal on the Shewhart charts and not on the CUSUM charts. However, if the T^2 statistic is sufficiently large, it will cause the subsequent S_n^+ statistics to be large, which may result in the CUSUM chart signaling a departure. In order to minimize the influence of any one T^2 statistic, especially in the initial time periods where near singular matrices result in large T^2 statistics, a Winsorizing constant (W) is used. The Winsorizing constant is the maximum allowable value that the T^2 statistic can take

when calculating the S_n^+ and S_n^- statistics for the multivariate CUSUM charts. When using a Winsorizing constant, the S_n^+ and S_n^- statistics are calculated as follows:

$$S_n^+ = \max(0, S_{n-1}^+ + \min(W, T_n^2) - k^+) \tag{4}$$

$$S_n^- = \min(0, S_{n-1}^- + \min(W, T_n^2) - k^-) \,. \tag{5}$$

This will prevent large values of T^2 from skewing the rest of the S_n^+ statistics in the CUSUM calculations and prevent the CUSUM charts from signaling a persistent shift. Winsorizing the T^2 statistic for the CUSUM charts will not effect the characteristics of the Shewhart charts. Shewhart chart will continue to use large un-Winsorized T^2 statistics to detect isolated departures in the data.

2. Database

The NATO Stabilization Force (SFOR) currently operating in Boznia-Herzegovina collects incident data on the local populace. This data is collected through numerous sources ranging from patrols of SFOR soldiers who personally encounter the local populace to theater level intelligence gathering sources. This data is divided into three categories based on the type of incident that occurred and the level of hostility contained in the act. The three categories are titled as follows: Threats and Rhetoric, Contentious Activities, and Violent Acts against SFOR. The data for each category is grouped into seven-day periods from Monday to Sunday in order to ensure significant data values in each category over each time period, to avoid confounding with the day of week, and to avoid sparseness.

The category "Threats and Rhetoric" is defined as acts of nonviolent demonstrations against SFOR, the international community or the local Boznia-Herzegovina government, as well as organized political statements against SFOR or the international community. Threats and Rhetoric contains such acts as radio broadcasts, peaceful demonstrations, and graffiti. Contentious Activities are defined as

acts that are controversial or suspicious in nature to either the international community or the Dayton Peace accord. Contentious Activities include such acts as demonstrations that hinder SFOR operations, observed vandalism of resettlement areas and material, confiscation of weapons by SFOR at weapon storage sites (WSS) or checkpoints, perceived acts of non-cooperation with established rules of the Dayton Peace accord by the local factions, and suspected intelligence gathering on SFOR units or bases by local nationals. Violence towards SFOR is defined as acts of outright violence towards SFOR personnel or facilities. Violence towards SFOR includes violent acts ranging from local personnel throwing rocks at SFOR patrols and vandalism against SFOR facilities to local personnel shooting at SFOR soldiers and acts of terrorism against SFOR personnel or facilities.

Even though the incident log received for this thesis was consolidated at the SFOR headquarters, units down to Battalion level maintain their own forms of incident logs for analysis. Military headquarters down to battalion level are staffed with personnel whose responsibility it is to consolidate and analyze enemy information. The incident logs at battalion level will normally not include incidents from outside their area of responsibility unless a higher headquarters has determined that a specific incident has implications for the lower units. The higher headquarters and lower units continuously exchange information in order to ensure that every level has a complete log of incidents and a complete understanding of the enemy situation. The SFOR incident log used in this thesis is listed in Appendix A.

3. Software

The software developed in this thesis is called "Multivariate CUSUM" and is an extension of the univariate CUSUM software package initially developed by Hawkins and Olwell and later modified by Weitzman. Multivariate CUSUM is in Microsoft Excel spreadsheet format and runs numerous macros in Visual Basic. The Microsoft Excel format

ensures its accessibility and usability to Army units down to battalion level, as well as most other organizations.

Multivariate CUSUM gives the user access to both univariate CUSUM procedures as well as the Multivariate CUSUM procedures developed in this thesis. From the main data worksheet, the user enters three data variates and then has the option of analyzing each variate individually or collectively. The main data page, "data1" is shown in Figure 13.

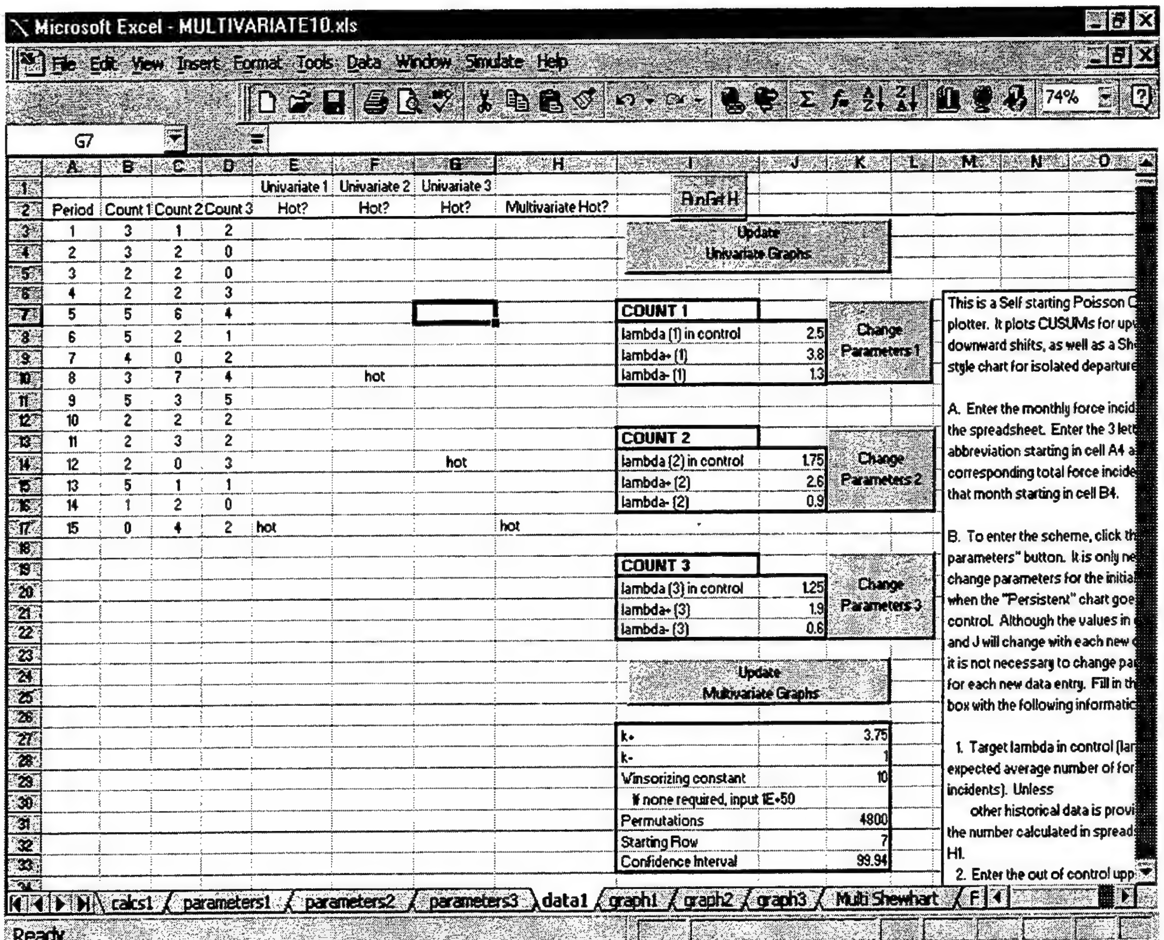

Figure 13. Multivariate Main Data Page, "data1". Column A is the time period entry field. Columns B, C, and D are the incident data entry fields. Columns E, F, G, and H are the out of control response fields for univariate and multivariate analysis respectively. The "Run Get H" button executes the ANYGETH.exe program. The "Update Univariate Graphs" button and the "Update Multivariate Graphs" button execute the respective programs and update the appropriate graphs. Change parameter buttons, which display a Visual Basic windows for entering CUSUM parameters (Figure 14), are shown for each variable along with the boxes used to calculate standard parameters as explained later in this section.

For univariate analysis, the user calculates the upper and lower CUSUM control chart limits for each individual variable using a Fortran based software package called "ANYGETH.exe" that was developed by Hawkins and Olwell. The user executes ANYGETH.exe by selecting the Visual Basic command button labeled "Run GET H". The user is prompted to input the proposed distribution of the data, and the in-control and out-of-control means. ANYGETH.exe returns the exact theoretical reference value k and prompts the user to input a reference value to use. Rounding the theoretical reference value k to the nearest .5 or .25 speeds the calculation of ANYGETH and yields satisfactory results. Next the user is prompted to input a Winsorizing constant, if necessary, and then to specify if he wants zero start or fast initial response (FIR) charts produced. Zero start charts are recommended and are used exclusively in this thesis. FIR charts are not used in this thesis, but are use to determine if the adjustments made to a restarted chart actually capture the nature of the shift that prompted the new chart. Finally the user is prompted to input the ARL. ANYGETH.exe returns multiple values of h and their corresponding ARL's.

For example, executing ANYGETH.exe and using a Poisson distribution with an in control mean of 3 and an out of control mean 5 returns an exact theoretical reference value of 3.915. Rounding this to 4 and using Zero start without a Winsorizing constant returns an upper control limit or decision interval (DI) of 6, and an in control ARL of 71.3. The user selects the DI for the upper control limit and inputs it into the excel worksheet. This process must be done separately for both the upward shift and the downward shift of each variable being analyzed.

Note that the exact desired ARL will often not be returned when using discrete data sets such as Poisson. The limited values of discrete data sets result in limited possible values of *h,* and also a limited set of possible ARL's (Hawkins and Olwell, 1998, p107-108).

The user inputs the parameters into the Excel program using the "Change Parameters" button from the main data page. This button opens another Visual Basic window, as shown in Figure 14, that prompts the user to input the persistent upper and lower control limits, the target Lambda in-control, Lambda+, Lambda-, and the isolated chart's probability limits.

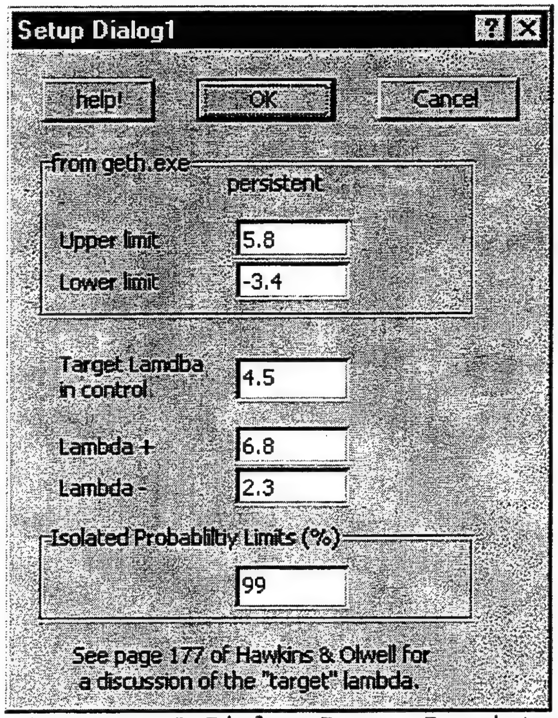

Figure 14. "Change Parameter" Dialog Box. Persistent upper and lower limits are values of the decision interval returned from ANYGETH.exe. Target Lambda in control, Lambda+, and Lambda- are parameters for which the CUSUM will be tuned to detect. The Isolated Probability Limits is the percentage used to calculate the initial Shewhart control limits.

The persistent upper and lower control limits are calculated using ANYGETH.exe. The target Lambda in-control, Lambda+, and Lambda- are determined by the commander or manager based off of the size of shift that he is concerned about. They may be calculated using the target mean of the variable times a constant or using a percentage of the target mean. In this thesis, the Lambda+ and Lambda- are calculated to detect a 50% shift in the target Lambda in-control. These values are automatically calculated on the main data page in the in cells designated for each data category.

The probability limits are used to calculate the initial upper and lower control limits for the isolated control charts and should be based off of the desired test ARL and equation 1 as explained in Chapter IV, section A2, above. Subsequent values for the upper and lower control limit are calculated using the CRITBINOM function explained earlier.

Once these parameters are entered, the user selects the "OK" button and returns to the main data page. The user executes the calculations and graphing by selecting the "Update Graphs" graphs. He is then able to view the graphs for each variable by selecting the appropriate worksheet sheet. Out of control signals will be shown both as "hot" values on the main data page and as points plotted outside the control limits on the graph pages.

It should be noted that the parameters only need to be changed when the charts have signaled a shift in data. The charts must then be cleared and the user will need to "retune" the charts to the new process mean.

For Multivariate analysis of the data, the user is able to input values for k^{+}, k^{-}, the Winsorizing constant, the confidence interval, the number of permutations, and the starting point into the main data page. For the number of permutations and the starting point, the values of 4800 and 7 respectively are suggested.

The user selects the "Update Multivariate Graphs" button, which executes the macro that conducts the nonparametric permutation technique described above in Chapter IV, Section C1. Conducting the nonparametric permutation technique for 1000 permutations may take considerable time if the data set is large. For example, on a Pentium III computer with a 300 mhz processor, 50 periods of data takes approximately 25 minutes to complete, and 100 periods of data takes nearly 90 minutes to complete. For this reason, the user is advised to make shorter runs when adjusting his values of k^{+} and k^{-}. When these parameters are adjusted, he can run the full 4800 permutations to ensure continuity of the control limits.

Multivariate CUSUM is designed for ease of use by personnel not highly trained in SPC and CUSUM techniques. It utilizes Microsoft Excel to ensure accessibility to a wide audience and Visual Basic Macro buttons to facilitate input of the required parameters. The general instructions for analyzing univariate and multivariate data, as described above, are displayed on the main data page. A copy of these instructions is located in Appendix B of this thesis.

V. STATISTICAL ANALYSIS

A. PARAMETER DETEMINATION

1. Individual Univariate Parameters

This thesis analyzes SFOR incident data from May 1999 to October 1999. In this section, we discuss the rationale and methods used to determine the numerous parameters required for individual univariate self-starting CUSUM control charts.

Individual univariate analysis consists of analyzing each data category individually using the univariate methods discussed previously. The control charts for a specific data category will only be restarted when a persistent shift is detected in that specific data category. The data categories are not combined with the other data categories, nor is the analysis of one data category dependent on the analysis conducted on the other data categories. This is not to be confused with simultaneous univariate analysis, which will be discussed in the next section.

In individual univariate analysis, the target in control mean (λ_o), or "Target Lambda in Control", is calculated by averaging the first four observations of the data set. For executing control charts with less than four observations, such as in the initial execution of the charts or when the charts are restarted as a result of a persistent shift, the target in control mean is calculated by averaging the available number of time periods, one through three. This follows the principal strength of self-starting CUSUM control charts, which is that they can be run with small initial data sets. Averaging larger amounts of data, such as seven or ten, increases the length of time required to determine the process mean and reduces the small data set strength of self-starting CUSUM control charts. The number of observations averaged is not related to the start up period of the multivariate control charts.

In the event that a CUSUM chart signals a shift and needs to be restarted, a new target in control mean must be calculated. The process is "retuned" by first looking at the graph and determining when the shift started, not when it was signaled. Shifts are said to start in the first time period following the time period where the trend line last touched the X-axis on the graph. The start point is also referred to as the first "shifted" data point because it is the first "shifted" data point following the last zero value of the trend line. The new in control mean is calculated in the same manner described above starting with the first "shifted" data point.

The upper and lower tuning parameters for the out of control means, λ_u and λ_1, or from the spreadsheet "Lambda+" and "Lambda-", are calculated using multiples of the in control mean. For this thesis, the out of control means are set to detect shifts of 50% of the actual mean. That is to say that Lambda+ is equal to three halves times the target sample mean and Lambda- is equal to one half times the target sample mean. In equation form: $\lambda^+ = 3/2*\lambda_o$ and $\lambda^- = 1/2*\lambda_o$. These values are used in order to detect large, "practically significant" shifts in the mean. Practically significant shifts refer to shifts in the mean that are deemed significant by the process manager. For example, managers that supervise the filling of oil tankers at port facilities use meters on their pumps to record the amount of oil pumped into a tanker ship. The tankers are subsequently charged for the amount of oil recorded by the meters. If the pumps or meters malfunction resulting in an average amount of 50 extra gallons of oil being pumped but not counted, the managers will probably not be concerned. The loss in revenue of these 50 gallons is insignificant to the total bill of loading a 5 million-gallon tanker. This is a "practically insignificant" shift and since charts will not be tuned to detect this shift, it will not be made "statistically significant".

If however, the limited capacity of the ship forces the extra 50 gallons of oil to be discarded into the ocean, and the pumping facility is fined $100,000 per spill, the over pumping will be a "practically significant" event. CUSUM and Shewhart charts will be tuned to detect this shift, making it "statistically significant."

The CUSUM chart upper and lower control limits (h^+ and h^-) are calculated using the Fortran software package "ANYGETH.exe". This software package requires the ARL and the univariate reference values (k^+ and k^-) to determine the upper and lower control limits.

This thesis chose a combined ARL of 100 for a number of reasons. First, the data is grouped into one-week periods running from Monday through Sunday. An ARL of 100 establishes the timeline of expecting a false alarm roughly once ever two years, which seemed reasonable. Secondly, in the area of military force protection, the cost of a false alarm is minimal compared to the cost of missing an upward shift, which warrants a low ARL. The cost of a false alarm includes increasing security measures and inconveniencing the soldiers when in fact the increase is unwarranted. The cost of missing a shift in the incident data may result in the loss of lives resulting from an incident such as the car bombing of the Air Force barracks, Khobar towers, in Saudi Arabia. Although the cost differences in this example are extreme, it is still favorable to avoid excessive false alarms. Besides inconveniencing the soldiers with increased force protection duties, excessive false alarms cause the soldiers to disregard the seriousness of their force protection duties. This sense of complacency degrades the effectiveness of the force protection and puts the soldiers at risk. An over all ARL of 100 is a compromise.

The individual univariate analysis uses four different tests per data category, which as stated in Chapter IV, requires special consideration in order to achieve the desired over all ARL. These four tests are the upper and lower Shewhart control limits, and the upper and

lower CUSUM control limits. From Equation 2, a test ARL of 400 is used in each of the four tests in order to obtain an overall process or combined ARL of 100.

As stated earlier, probability limits are used to determine the values of the upper and lower Shewhart control limits for the first data point. Control limits for subsequent data points are calculated by the CRITBINOM function. From Equation 1, probability limits of .9975, or 99.75%, result in the desired test ARL of 400.

The initial univariate reference values, k^+ and k^-, and the upper and lower control limits for the SFOR data were determined using the previously mentioned software package "ANYGETH.exe". The results of this work are consolidated in table 1 below.

Data Category	k+/k-	h+/h- (DI)	In Control ARL	Out of Control ARL
1 Threats and Rhetoric	8.6 (+) 5 (-)	10.8 (+) -7 (-)	417 up 469 down	6.3 up 5 down
2 Contentious Activities	11.4 (+) 6.7 (-)	10.8 (+) -6.6 (-)	404 up 411 down	5 up 3.9 down
3 Violence Towards SFOR	3.1 (+) 1.8 (-)	9.3 (+) -6.2 (-)	418 up 430 down	13 up 11.9 down

Table 1. Results of ANYGETH.exe on SFOR data. Winsorizing constant was not used. Up corresponds to upward shifts, down corresponds to downward shifts.

2. Multivariate Parameters

As stated earlier, multivariate analysis consists of two parts: simultaneous univariate analysis and nonparametric multivariate analysis. The simultaneous univariate analysis parameters are calculated in the same manner as the individual univariate analysis parameters. One difference is that multivariate analysis has 16 tests, twelve in the simultaneous univariate analysis and four in the nonparametric multivariate analysis, which affect the combined ARL. Using Equation 2, a desired test ARL of 1600 will achieve the combined ARL of 100. This test ARL of 1600 is used for the 12 simultaneous univariate tests. The test ARL of 1600 also affects the probability

limits used for the Shewhart control limits of the chart's first data point. Using Equation 1, a probability limit of .999375, rounded to 99.94%, achieves an in control ARL of 1667, which is sufficiently close to the desired in control ARL of 1600.

The nonparametric multivariate analysis developed in this thesis requires only four principal parameters: the multivariate reference values k^+ and k^-, the confidence interval, and a Winsorizing constant. The reference values k^+ and k^- were determined by running multiple simulations on the data using different values and determining which values resulted in the flattest control limits. The initial values of k^+ and k^- were set to 4 and 2, but after several simulations on the SFOR data set, the values were changed to 3.75 and 1 for reasons described earlier.

The same methodology was used to determine the value of the Winsorizing constant. After running several simulations with different Winsorizing constants, this thesis chose a Winsorizing constant of 10 because it limited the effect extreme values of T^2 had on the values of S_n^+ and S_n^-.

As with the probability limits in the simultaneous univariate analysis, the confidence interval chosen for the control limits in the nonparametric permutation technique directly affects the in control ARL. Again from Equation 1, a confidence interval of .999375, rounded to 99.94%, achieves an in control ARL of 1667, which is sufficiently close to the desired in control ARL of 1600. This nonparametric multivariate test ARL, when combined with the simultaneous univariate test ARL of 1600 using Equation 2, results in an over all ARL of 101.015, which is acceptably close to the target combined ARL of 100. The out of control ARL will not be discussed in the multivariate analysis. This is due to the fact that the out of control ARL depends on the type of shift that occurs. In multivariate analysis, numerous types of shifts can occur.

Attempting to address all possible shifts, or even focus on a few, is beyond the scope of this thesis. As a result, we also do not discuss power considerations.

As stated earlier, the user may choose to change the number of permutations and the start point of the nonparametric permutation technique. Manipulating the number of permutations and the start point are not self-explanatory and require further explanation.

Manipulating the number of permutations affects the time of the program operation, the smoothness of the control limits, and the thoroughness of the sampling. It should, however, be based on the confidence interval used to get the proper multivariate test ARL. Obviously, the fewer the permutations, the quicker the program executes the technique. But this also increases the variance in the estimates of the control limits and should leave the user less confident that the control limits reflect the correct percentile of possible values from the sample. Also, if high ARL's are used, a high number of permutations should be used to prevent the control limits from taking on the extreme points of the permutated values. For example, using a confidence interval of 99.94% on 100 permutations of the data will result in the highest and lowest values of the permutated statistics. On the other hand, using 50,000 permutations will result in the 49,970th and 30th sorted values of the permutated statistic for the upper and lower control limits. This additional distance from the highest and lowest values provides additional confidence that the control limits are not affected by extreme values. Of course time and computing power will effect the final decision as well. This thesis chose to conduct 4,800 permutations on the data making the 4797th and 3rd sorted values of the permutated statistic the upper and lower control limits.

Manipulating the start point for the calculations will effect the initial values of the T^2 statistic. If the start point is equal to the number of variables, near singular covariance matrices are common.

These near singular covariance matrices will cause the T^2 statistic to take an extreme value, which in turn will skew the Shewhart style graphs. Using a start point equal to three or four time periods past the number of variables produced large, but not extreme values of T^2. Through simulation, this thesis determined that a start point of 7 was acceptable, in that it reduced the start up time for the graphs while producing usable values of T^2.

B. APPLICATION TO STABILIZATION FORCE (SFOR) DATA

1. Individual Univariate Analysis

This thesis will conduct individual univariate analysis on all three data categories, but will only discuss the results of the first category in detail. The results of the analysis on the second and third data categories will be consolidated at the end of this section. Multivariate analysis of the data, consisting of simultaneous univariate analysis and nonparametric multivariate analysis, will be conducted and discussed in the following section.

The individual univariate control charts for data category 1, Threats and Rhetoric, are shown below in Figure 15. The parameters used in the charts are those listed in Table 1.

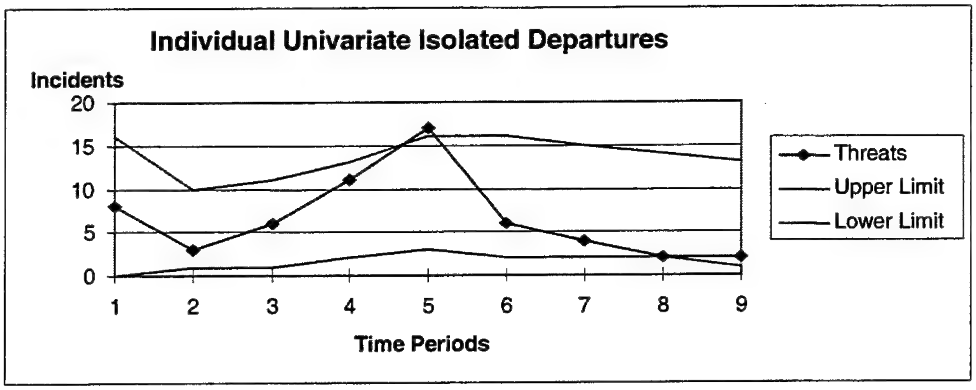

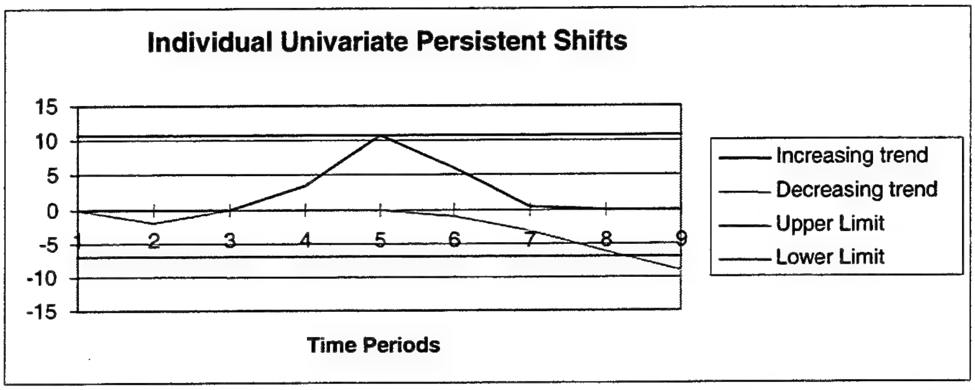

Figure 15. Individual Univariate Control Charts for SFOR Data, Threats and Rhetoric, Periods 1-9. Isolated upward departure at time period 5 and a persistent downward shift at time period 9. The persistent decreasing shift appears to begin at time period 6.

These charts signaled an isolated upward departure at time period 5 and a persistent downward shift at time period 9. Although close, the increasing trend line on the persistent chart does not exceed the upper control limit at time period 5 and therefore, does not signal a shift. This can be verified in Excel by selecting the increasing trend line with the pointer arrow. When the pointer arrow is placed on the selected trend line near the point corresponding to time period 5, Excel displays the value of the increasing trend line at time period 5 as 10.735. This is less than the upper control limit of 10.8 and a persistent shift is not signaled.

The charts need to be retuned for the persistent shift, not for the isolated departure. The charts are restarted at the point where the shift started, not when it signaled. The start of a shift is identified by the first time period following the time period where the trend line

last touched the X axis on the graph. The start point is also referred to as the first "shifted" data point because it is the first "shifted" data point following the last zero value of the trend line. From Figure 15, the persistent downward shift detected at time period 9 was last plotted on the X-axis at time period 5. The next point after that, or the first "shifted point", is at time period 6. The new charts are therefore retuned and restarted at time period 6.

Figure 16 shows the updated charts, started at time period 6, that are tuned to detect shifts from the new process mean. The new target in control mean is 3.5 which is a considerable decrease in the target in control mean from previous in control mean. The new out of control mean for an upward shift is 5.3, and the new out of control mean for a downward shift is 1.8. The upper and lower control limits are 10 and -7 respectively. The ARL is 413 for the upward shift and 411 for the downward shift.

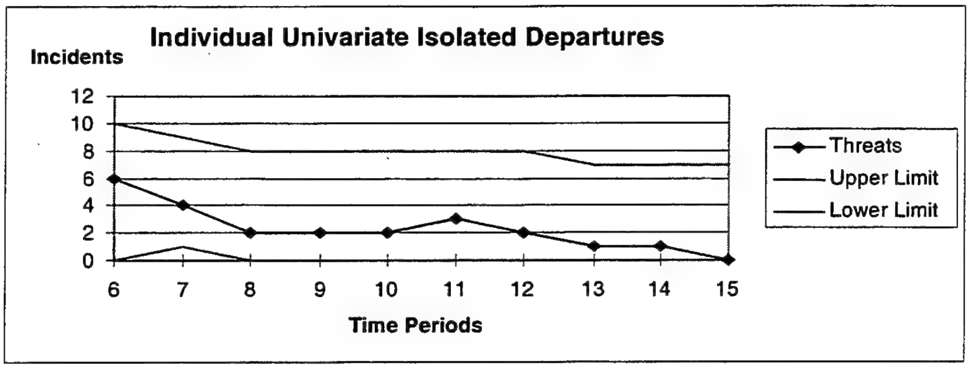

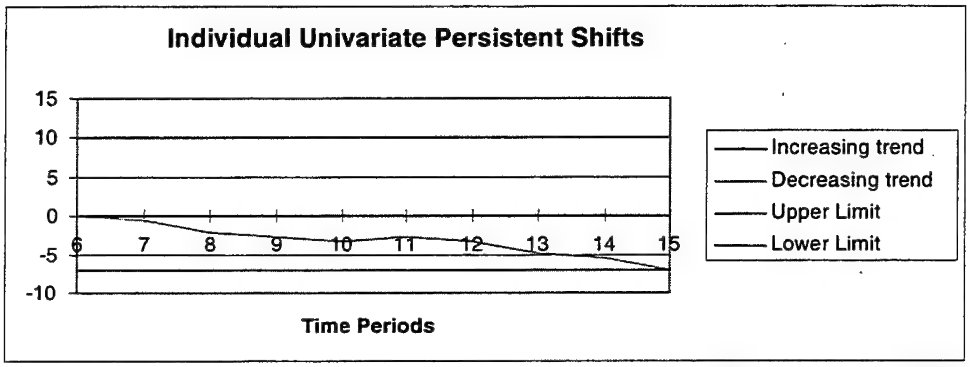

Figure 16. Individual Univariate Control Charts for SFOR Data, Threats and Rhetoric, Periods 6-15. Persistent downward shift signaled at time period 15. Decreasing trend appears to begin at time period 7.

The charts in Figure 16 signal persistent downward shift at time period 15, which appears to start at time period 7. The fact that the shift appears to start immediately following the start period of the newly tuned charts suggests that the shift was not the result of a step change, but is instead the result of a linear drift in the data. When retuning and restarting a chart due to a shift caused by linear drift, the chart is restarted at the first time period after the shift was detected. In this case, the new chart will start at time period 16.

Restarting the CUSUM charts at time period 16, however, illustrates the issue of starting a CUSUM chart with an initial value equal to zero. CUSUM charts require an initial value not equal to zero. If they are started with an initial value equal to zero, the charts will signal a persistent shift in the time period that contains the first non-zero value. This issue presented itself throughout the analysis of SFOR incident data due to the number of time periods that contain values

equal zero. To avoid this issue, this thesis will restart the charts in the first non-zero time period after the apparent start of the shift. In this case, time periods 16 and 17 contain zero values, so the charts will be started in time period 18.

Figure 17 shows the updated charts that are tuned to detect shifts from the new process mean. The new target in control mean is 0.25, which is another decrease in the target in control mean from the previous in control mean. The new out of control mean for an upward shift is 0.4, and the new out of control mean for a downward shift of 0.1. The upper and lower control limits are 6.1 and -3.6 respectively. The in control ARL for the upward shift is 409 and the in control ARL for the downward shift is 412.

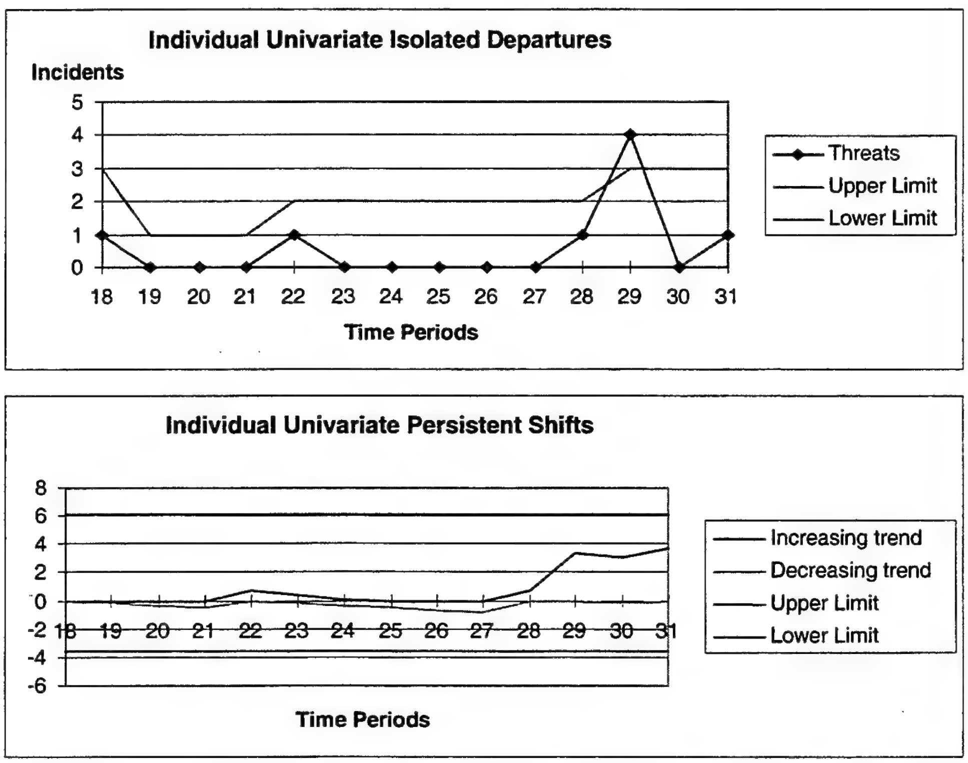

Figure 17. Individual Univariate Control Charts for SFOR Data, Threats and Rhetoric, Periods 18-31. Isolated upward departure at time period 29. Process is in control.

The new charts in Figure 17 detected an isolated upward departure at time period 29. There were no persistent shifts detected, therefore

the system is in control through time period 31, which is the end of the observed data.

Table 2 below shows the consolidated results of the univariate analysis for the 3 data categories.

			INDIVIDUAL UNIVARIATE ANALYSIS							
Data Category	**Time Periods**	**Target In Control Mean**	**Out of Control Mean**	**k+/k-**	**h+/h- (DI)**	**In Control ARL**	**Out of Control ARL**	**Isolated Departures**	**Persistent Shifts**	**Type of Persistent Shift**
1 Threats & Rhetoric	1-9	7	10.5 up 3.5 down	8.6 (+) 5 (-)	10.8 (+) -7 (-)	417 up 469 down	6.3 up 5 down	up at 5	down at 9	step
	6-15	3.5	5.3 up 1.8 down	4.3 (+) 2.6 (-)	10 (+) -7 (-)	413 up 411 down	10.3 up 8.7 down	n/a	down at 15	linear drift
	18-31	0.25	.4 up .1 down	.3 (+) .16 (-)	6.1 (+) -3.6 (-)	409 up 412 down	51.1 up 44.5 down	up at 29	n/a	n/a
2 Contentious Activities	1-16	9.25	13.9 up 4.6 down	11.4 (+) 6.7 (-)	10.8 (+) -6.6 (-)	404 up 411 down	5 up 3.9 down	up at 14	down at 16	step
	15-31	3.25	4.9 up 1.6 down	4 (+) 2.3 (-)	11 (+) -6 (-)	569 up 403 down	11.9 up 8.7 down	n/a	n/a	n/a
3 Violence Toward	1-18	2.5	3.8 up 1.3 down	3.1 (+) 1.8 (-)	9.3 (+) -6.2 (-)	410 up 414 down	13 up 11.9 down	up at 4	down at 18	step
	11-22	1	1.5 up .5 down	1.2 (+) .7 (-)	8.8 (+) -5.2 (-)	418 up 430 up	26.5 up 22.4 down	n/a	n/a	n/a

Table 2. Consolidated Individual Univariate Analysis on SFOR Incident Data. Up corresponds to upward shifts and down corresponds to downward shifts.

From Table 2, the number of shifts in the three categories suggests high volatility in the SFOR incident data and of the peacekeeping environment itself. Using test ARL's of 400 in the four combined tests for each data category should have resulted in one false alarm every 100 time periods. Instead, each data category had at least one shift in only 31 time periods. This is three times as many shifts as would be expected and clearly shows the volatility of the situation.

2. Multivariate Analysis

The initial parameters for the simultaneous univariate analysis and the nonparametric multivariate analysis are listed below in Table 3. The simultaneous univariate parameters are entered and the corresponding charts are updated. Following this, the multivariate parameters are entered and the nonparametric permutation technique is conducted. All charts are restarted simultaneously if a persistent shift is detected in any of the CUSUM control charts.

As described earlier, the multivariate parameters are the reference values (k^+ *and* k^-), the Winsorizing constant, and the

confidence interval. These four parameters are set at 3.75, 1, 10, and 99.94% respectively and will remain so throughout the multivariate analysis unless a change is required. The additional parameters, the number of permutations and the start point, are set at 4800 and 7 respectively. These parameters will also remain constant throughout the multivariate analysis unless a change is required.

Executing the simultaneous univariate analysis resulted in two isolated departures and one persistent shift as shown below in Table 3. The first persistent shift in multivariate analysis occurs as a univariate persistent downward shift in category 1, Threats and Rhetoric, in period 9 and it appears to start at time period 6. There was also an isolated upward departure at time period 5. There were no shifts detected in the nonparametric multivariate control charts.

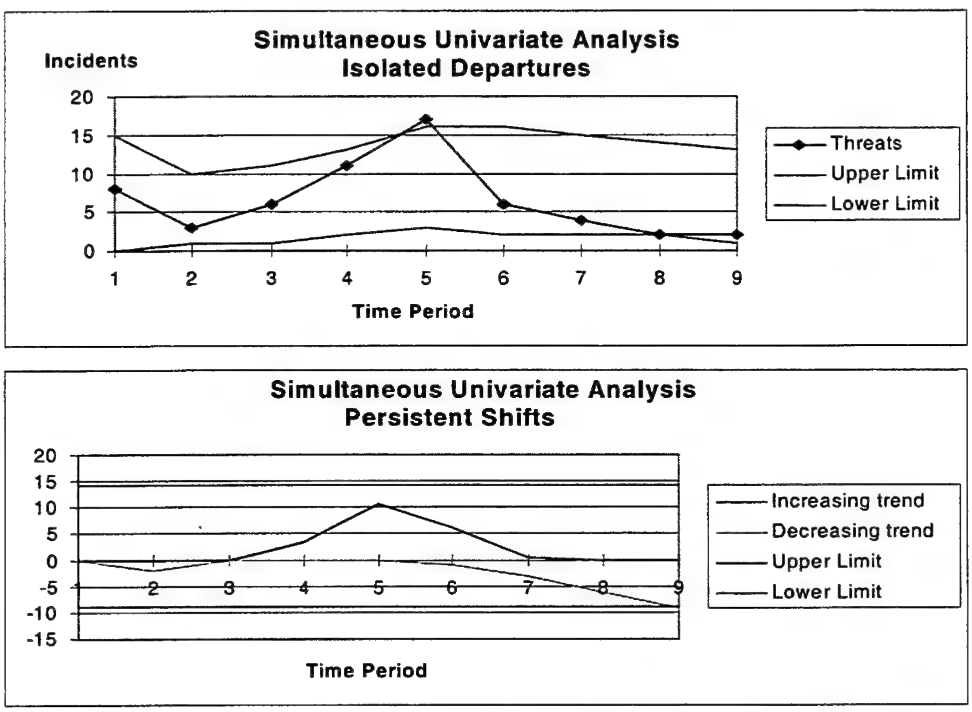

Figure 18. Simultaneous Univariate Analysis, Persistent Shift in Threats and Rhetoric, Time Periods 1-9. Isolated upward departure at time period 5. Persistent downward shift at time period 9. Persistent downward shift appears to start at time period 6.

The parameters and results of the analysis are consolidated in Table 3 below.

Data Category	Time Periods	Target In Control Mean	Out of Control Mean	k+/k-	h+/h- (DI)	In Control ARL	Out of Control ARL	Isolated Departures	Persistent Shifts	Type of Persistent Shift
			SIMULTANEOUS UNIVARIATE PARAMETERS							
1 Threats & Rhetoric	1-9	7	10.5 up 3.5 down	8.6 (+) 5 (-)	14.2 (+) -9 (-)	1646 up 1985 down	8.1 up 6.3 down	up at 5	down at 9	step
2 Contentious Activities	1-9	9.25	13.9 up 4.6 down	11.4 (+) 6.7 (-)	14.2 (+) -8.9 (-)	1643 up 1860	6.4 up 4.9 down	n/a	n/a	n/a
3 Violence Toward	1-9	2.5	3.8 up 1.3 down	3 (+) 1.8 (-)	15 (+) -8.2 (-)	1980 up 1693 down	18.4 up 15.8 down	up at 4	n/a	n/a

k+	k-	Confidence Interval	Winsorizing Constant	Iterations	Start Point	Isolated Departures	Persistent Shifts	Type of Persistent Shift
		NONPARAMETRIC MULTIVARIATE PARAMETERS						
3.75	1	99.94%	10	4800	7	n/a	n/a	n/a

Table 3. Consolidated Parameters, Multivariate Analysis, Time Periods 1-9. Up corresponds to upward shifts and down corresponds to downward shifts. Persistent downward shift detected in the simultaneous univariate control charts in data category 1 at time period 9. No shifts detected in the nonparametric multivariate control charts.

The persistent shifts require that all the charts be restarted. All charts, both the simultaneous univariate charts and the nonparametric multivariate charts, will be restarted using the first detected shift. All categories will be restarted at this time even though there has not been a signaled shift in a multivariate chart. Since the first persistent shift appears to start at time period 6, the new charts will be restarted at time period 6.

The consolidated parameters and results of the analysis for time periods 6-21 are shown below in Table 4.

Data Category	Time Periods	Target In Control Mean	Out of Control Mean	k+/k-	h+/h- (DI)	In Control ARL	Out of Control ARL	Isolated Departures	Persistent Shifts	Type of Persistent Shift
			SIMULTANEOUS UNIVARIATE PARAMETERS							
1 Threats & Rhetoric	6-21	3.5	5.3 up 1.8 down	4 (+) 1.8 (-)	19 (+) -8.2 (-)	1755 up 1693 down	15 up 15.8 down	n/a	n/a	n/a
2 Contentious Activities	6-21	3.75	5.6 up 1.9 down	4.6 (+) 2.7 (-)	13.6 (+) -8.3 (-)	1629 up 1606 down	13.7 up 10.6 down	n/a	down at 21	step
3 Violence Toward	6-21	2.25	3.4 up 1.1 down	2.8 (+) 1.6 (-)	12.4 (+) -8 (-)	1733 up 1852 down	19.9 up 15.6 down	n/a	n/a	n/a

k+	k-	Confidence Interval	Winsorizing Constant	Iterations	Start Point	Isolated Departures	Persistent Shifts	Type of Persistent Shift
		NONPARAMETRIC MULTIVARIATE PARAMETERS						
3.75	1	99.94%	10	4800	7	n/a	n/a	n/a

Table 4. Consolidated Parameters, Multivariate Analysis, Time Periods 6-21. Up corresponds to upward shifts and down corresponds to downward shifts. Persistent downward shift detected in the simultaneous univariate control charts in data category 2 at time period 21. No shifts detected in the nonparametric multivariate control charts.

As can be seen in Table 4, the only shift occurred in category 2, Contentious Activities. It is a persistent downward shift detected in the simultaneous univariate CUSUM control chart. No shifts are detected with the nonparametric multivariate control charts. The shift is the result of a step change and appears to start at time period 14, so the new charts will be restarted at time period 14.

Restarting the CUSUM charts time period 14, again illustrates the issue of starting a CUSUM chart with an initial value equal to zero. As stated earlier, CUSUM charts require an initial value not equal to zero. In the event of a zero value in an initial chart time period, this thesis stated earlier that it would start the charts at the next time period with a non-zero value.

Time periods 14-21 contain zeros in one category or the other. Restarting the charts at time period 22 would result in the loss of eight time periods, or 2 months worth of data. To prevent the loss of such a significant amount of data, the original rule will be broken and the charts will be started at time period 13, which is the first time period prior to the start of the shift with all non-zero values. The parameters used and the results of the analysis are consolidated in Table 5.

				UNIVARIATE PARAMETERS						
Data Category	Time Periods	Target In Control Mean	Out of Control Mean	k+/k-	h+/h- (DI)	In Control ARL	Out of Control ARL	Isolated Departures	Persistent Shifts	Type of Persistent Shift
1 Threats & Rhetoric	22-31	0.25	.4 up .1 down	.3 (+) .2 (-)	9.3 (+) -7.8 (-)	1607 up 1759 down	82.8 up 73.5 down	up at 29	n/a	n/a
2 Contentious Activities	22-31	2.25	3.4 up 1.1 down	2.8 (+) 1.6 (-)	12.4 (+) -8 (-)	1733 up 1852 down	19.9 up 15.6 down	n/a	n/a	n/a
3 Violence Toward	22-31	0.75	1.1 up .4 down	.9 (+) .6 (-)	11.7 (+) -9.6 (-)	1629 up 1682 down	52.3 up 44.8 down	n/a	n/a	n/a

			MULTIVARIATE PARAMETERS					
k+	k-	Confidence Interval	Winsorizing Constant	Iterations	Start Point	Isolated Departures	Persistent Shifts	Type of Persistent Shift
3.75	1	99.94%	10	4800	7	n/a	n/a	n/a

Table 5. Consolidated Parameters, Multivariate Analysis, Time Periods 13-31. Up corresponds to upward shifts and down corresponds to downward shifts.

There were two isolated departures detected during time periods 13-31. One shift was an isolated upward departure at time period 29 in category 1, Threats and Rhetoric. The other isolated departure was an isolated downward departure at time period 14 in category 2, Contentious Activities. The charts do not need to be restarted since there were no persistent shifts detected. The process is in control through the end of the data set.

The shifts that occurred during the multivariate analysis were all from the simultaneous univariate charts. Figures 19, 20, and 21 consolidate all these departures and shifts on one graph per data category. Large red data points identify the detected shifts and departures.

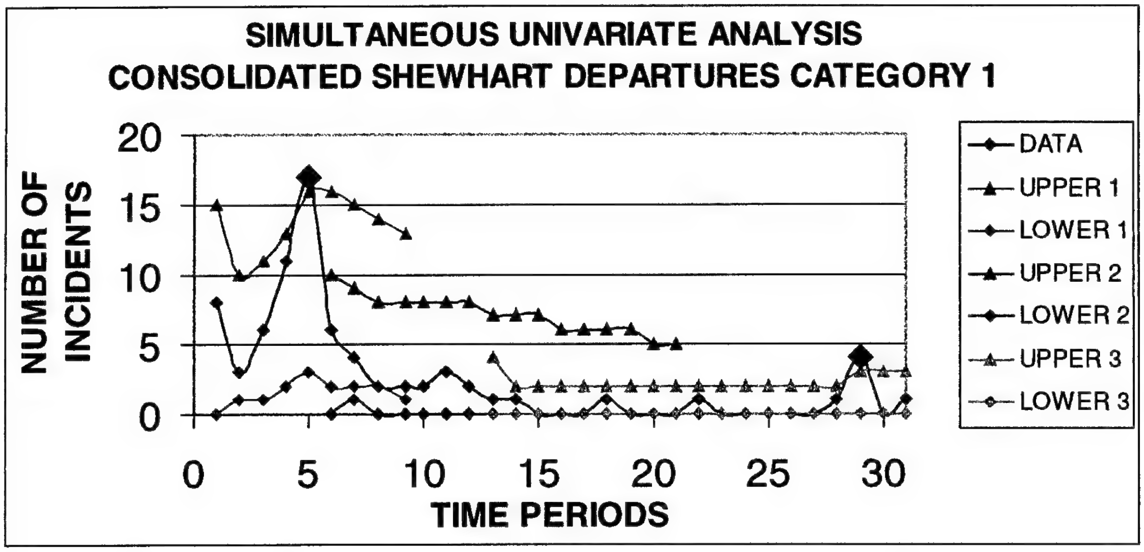

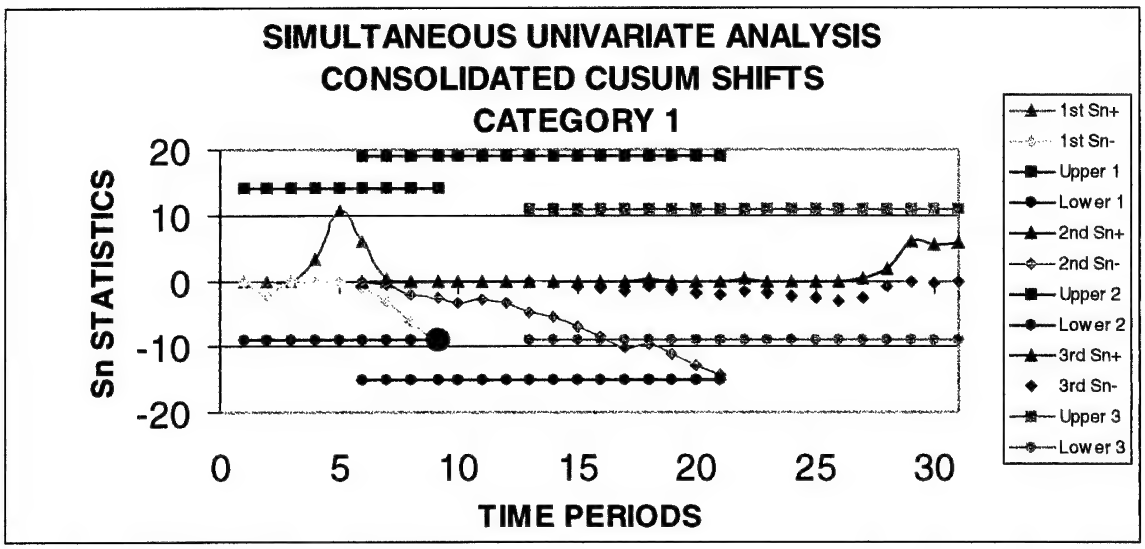

Figure 19. Simultaneous Univariate Analysis, Consolidated Shifts in Category 1. 1st chart periods are from time period 1 to time period 9. 2nd chart periods are from time period 6 to time period 21. 3rd chart periods are from time period 13 to time period 31. Isolated departures were detected in time periods 5 and 29. One persistent shift occurred in time period 9. Large red data point identifies shifts and departures.

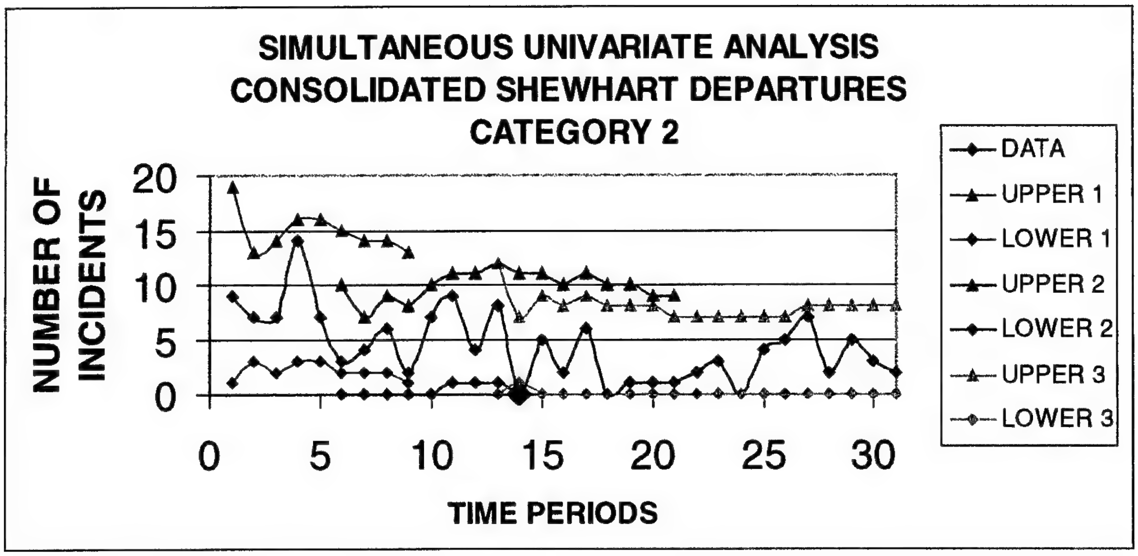

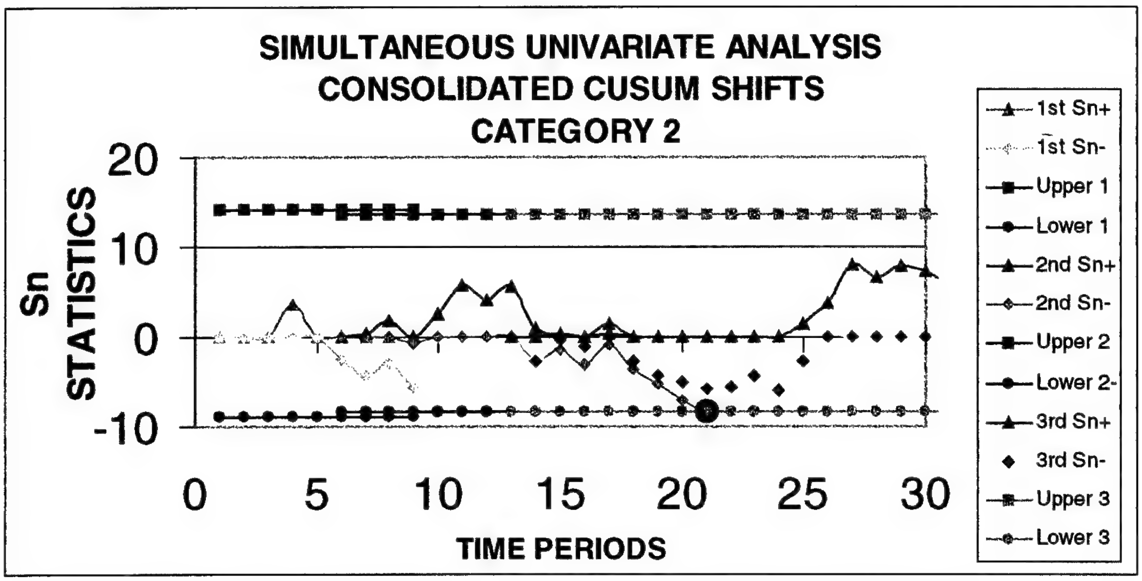

Figure 20. Simultaneous Univariate Analysis, Consolidated Shifts in Category 2. 1st chart periods are from time period 1 to time period 9. 2nd chart periods are from time period 6 to time period 21. 3rd chart periods are from time period 13 to time period 31. An isolated departure occurred in time period 14. A persistent shift occurred in time period 21. Large red data point identifies shifts and departures.

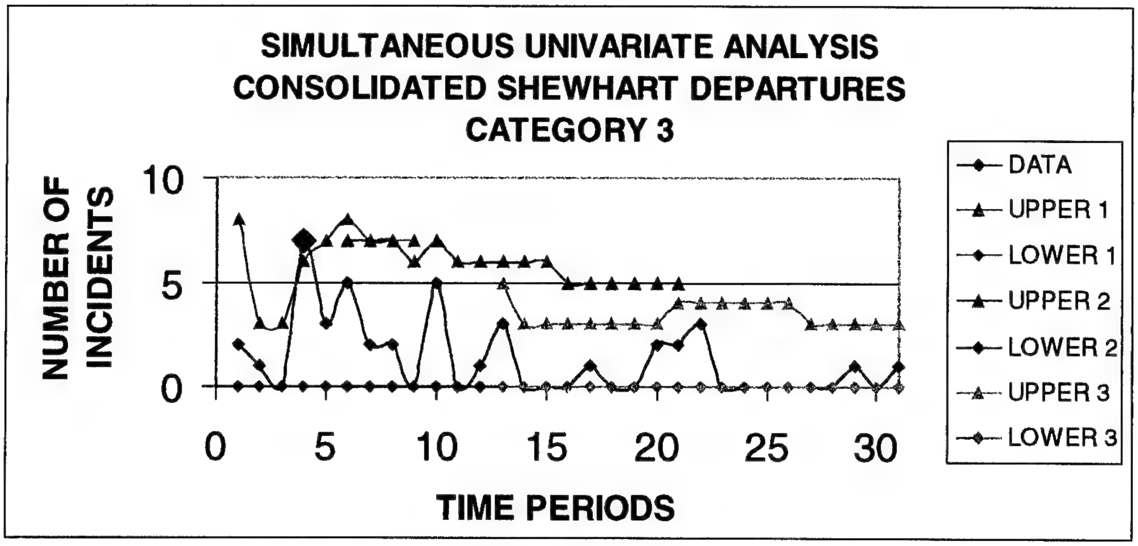

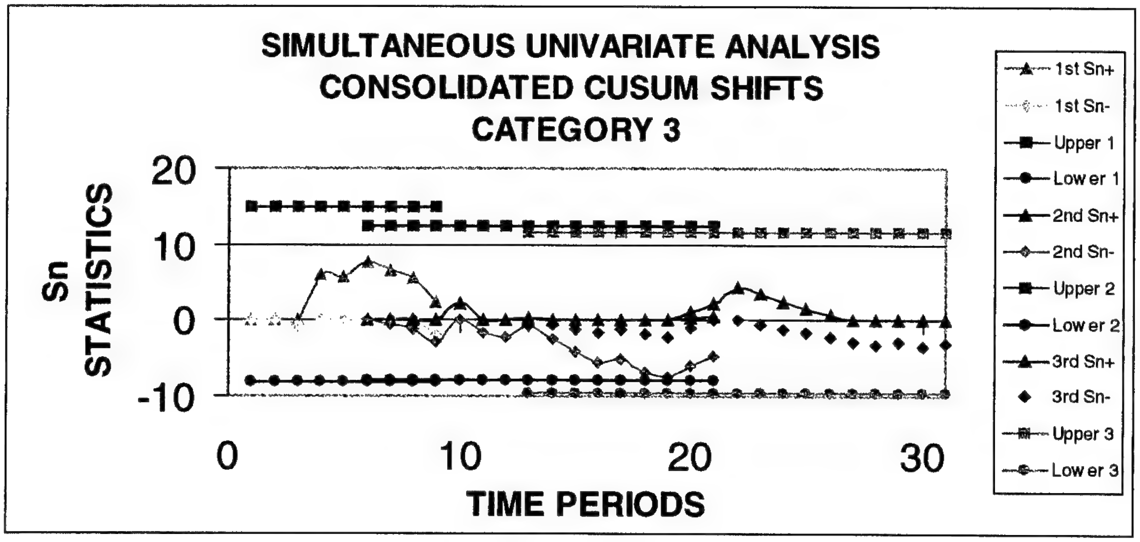

Figure 21. Simultaneous Univariate Analysis, Consolidated Shifts in Category 3. 1st chart periods are from time period 1 to time period 9. 2nd chart periods are from time period 6 to time period 21. 3rd chart periods are from time period 13 to time period 31. Isolated departure occurred in time period 4. No persistent shifts were detected. Large red data point identifies departure.

3. Analysis of SFOR Incident Data in Reverse Order

Applying the technique developed to the actual SFOR data, as done above, shows volatile data with primarily decreasing trends. To a commander responsible for the lives of his soldiers, decreasing trends which warrant a decrease in the force protection level do not stimulate the same sense of anxiety as increasing trends would. Obviously, increasing trends depict a situation that is getting worse, and for the

commander, a situation where his soldiers are in significantly more danger.

To show the results of this technique on increasing trends, this thesis reversed the order of the SFOR incident data, then applied these techniques to it. The numbers of incidents should now be generally increasing instead of decreasing, which will signal more increasing trends. Again, we will analyze the data category 1, Threats and Rhetoric, in detail and summarize the individual univariate analysis of data categories 2 and 3. Following the individual univariate analysis, we will analyze the data using the multivariate technique.

Starting with the individual univariate analysis of category 1, the reversed data has a target in control mean of 1.5, an out of control mean for an upward shift of 2.3, and an out of control mean for a downward shift of .8. The upper control limit equals 8 and the lower control limit equals -6. The in control ARL for an upward shift is 404 and the in control ARL for a downward shift is 403. The results are shown below in Figure 22.

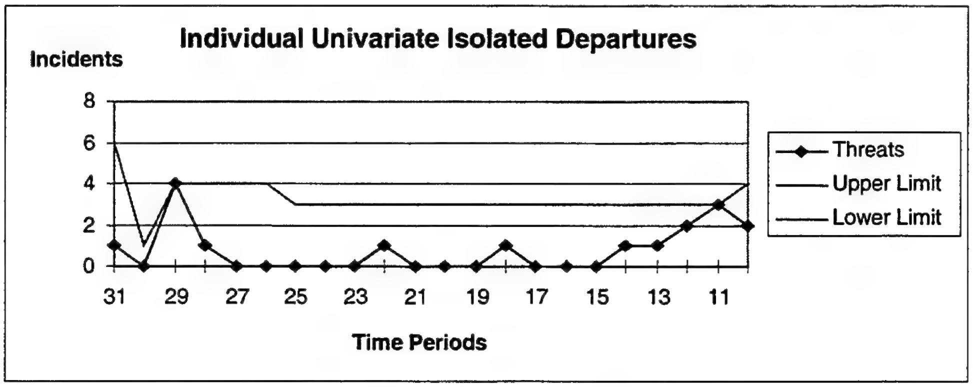

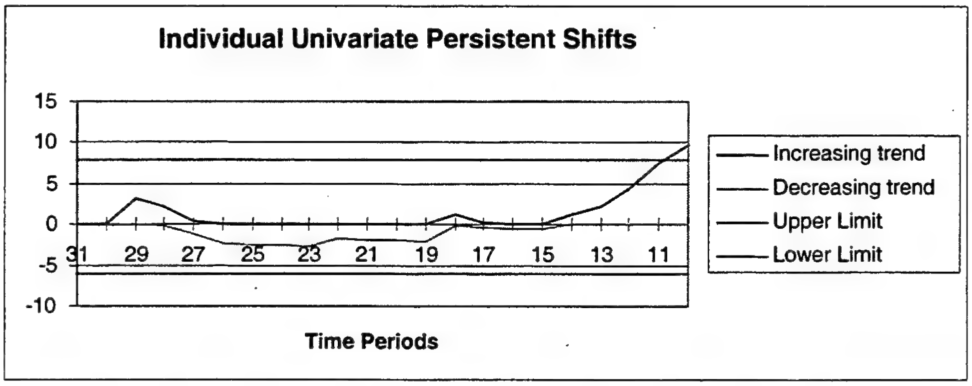

Figure 22. Individual Univariate Control Charts for Reversed SFOR Data, Threats and Rhetoric, Periods 31-10. Isolated upward departures at time periods 29 and 11. Persistent upward shift at time period 10. Persistent shift appears to begin at time period 14.

As shown, two isolated upward departures are detected at time periods 29 and 11. A persistent upward shift is detected at time period 10, which appears to start at time period 14. This is a step change. The charts need to be retuned and restarted at time period 14.

The new charts for time periods 14 through 5 are shown below in Figure 23. The new target in control mean is 1.75, the out of control mean for an upward shift is 2.6, and the out of control mean for a downward shift is .9. The upper control limit is equal to 9.7 and the lower control limit is equal to -6.4. The in control ARL for an upward shift is 404 and the in control ARL for a downward shift is 403.

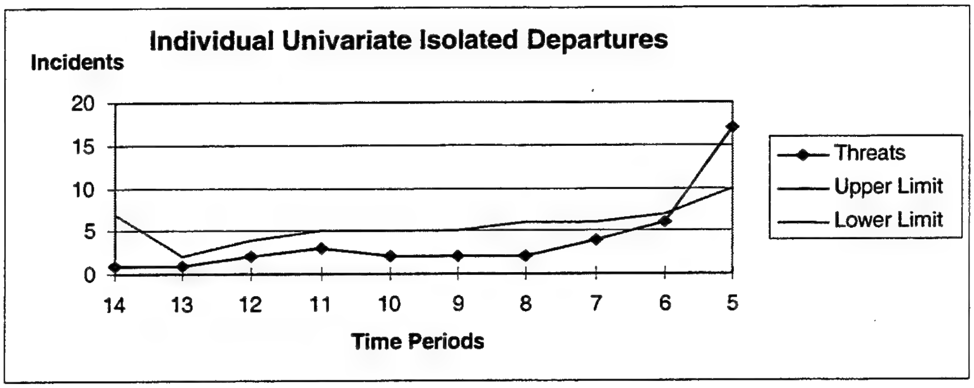

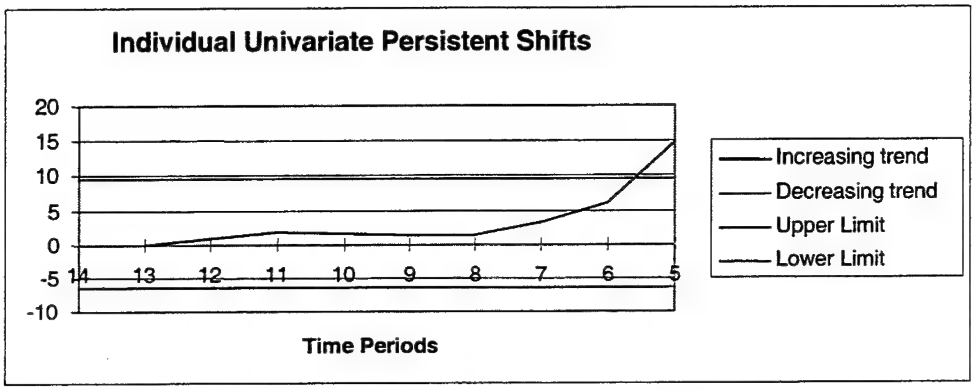

Figure 23. Individual Univariate Control Charts for Reversed SFOR Data, Threats and Rhetoric, Periods 14-5. Isolated upward departure and persistent upward shift at time period 5. The persistent upward shift appears to begin at time period 12.

Time period 14 through 5 show an isolated upward departure and a persistent upward shift at time period 5. The persistent shift appears to start at time period 12. Once again, this is a step change and the new charts will be restarted at time period 12.

The new charts restarted at time period 12 have a target in control mean of 2.25, an out of control mean for an upward shift of 3.4, and an out of control mean for a downward shift of 1.1. The upper and lower control limits are 9.2 and -6 respectively. The in control ARL's are 433 for an upward shift and 419 for a downward shift. The results are shown below in Figure 24.

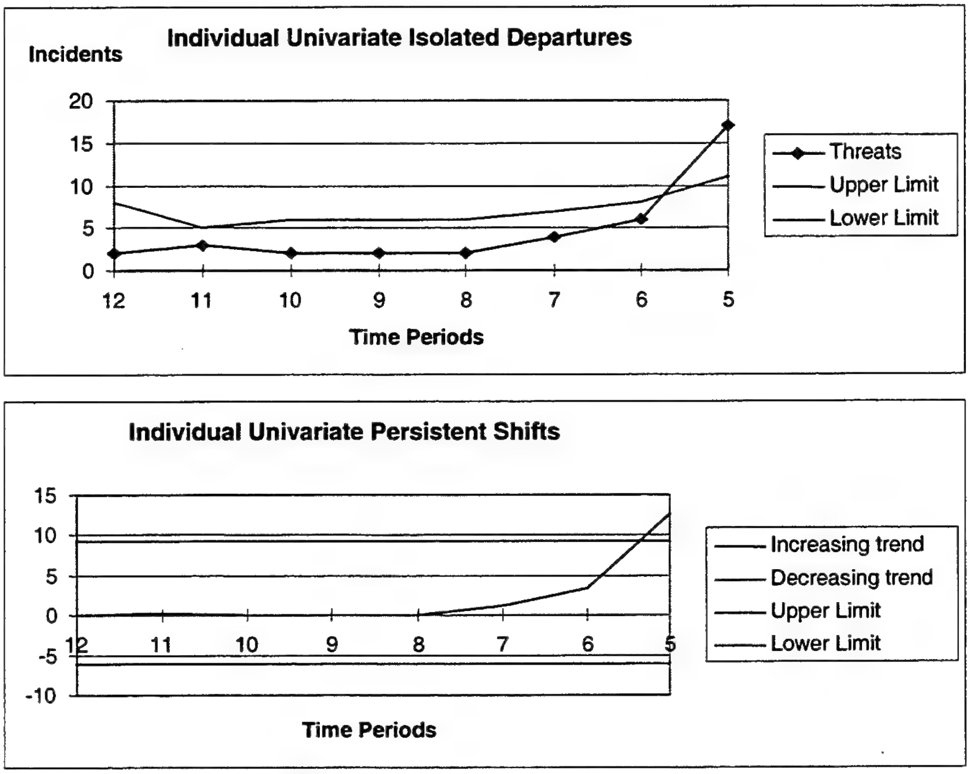

Figure 24. Individual Univariate Control Charts for Reversed SFOR Data, Threats and Rhetoric, Periods 12-5. Isolated upward departure and persistent upward shift at time period 5. The persistent upward shift appears to begin at time period 7.

The charts signal once again signal an isolated departure and a persistent shift at time period 5. The persistent shift appears to start at time period 7, depicting a step change. The charts will be restarted at time period 7.

Figure 25 shows the restarted charts for time periods 7 through 1. The new target in control mean is 9.5, the out of control mean for an upward shift is 14.3, and the out of control mean for a downward shift is 4.8. The upper and lower control limits are 10.7 and -6.8 respectively. The in control ARL's are 406 for an upward shift and 414 for a downward shift.

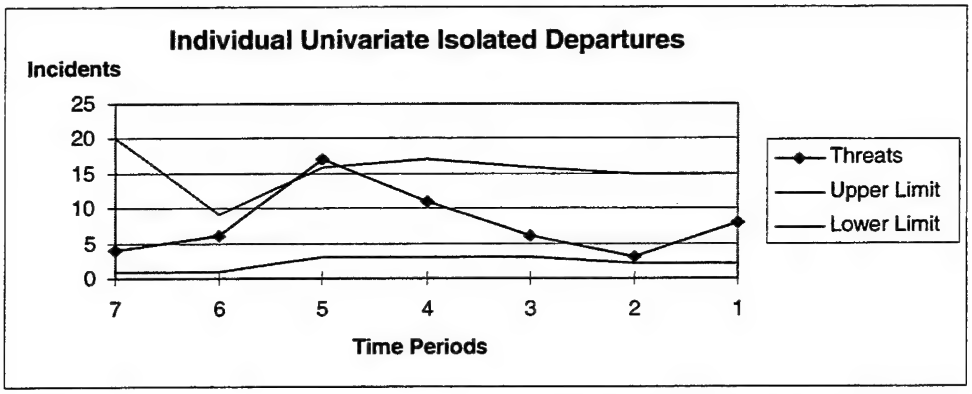

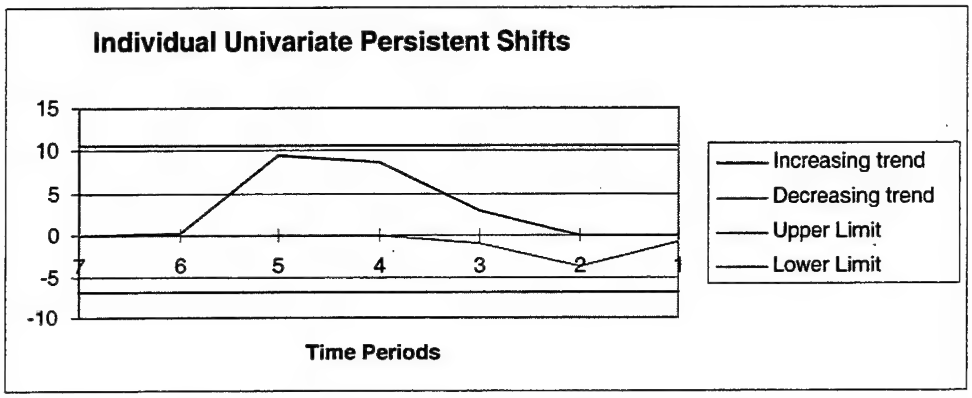

Figure 25. Individual Univariate Control Charts for Reversed SFOR Data, Threats and Rhetoric, Periods 7-1. Isolated upward shift signaled at time period 5.

The charts detect an isolated upward departure at time period 5. There were no persistent shifts detected so the process is in control.

The consolidated results from the univariate analysis of the SFOR data in reverse order is shown below in Table 6.

			INDIVIDUAL UNIVARIATE ANALYSIS							
Data Category	Time Periods	Target In Control Mean	Out of Control Mean	k+/k-	h+/h- (DI)	In Control ARL	Out of Control ARL	Isolated Departures	Persistent Shifts	Type of Persistent Shift
1 Threats & Rhetoric	31-10	1.5	2.3 up .8 down	1.9 (+) 1.1 (-)	8 (+) -6 (-)	404 up 403 down	18.1 up 18.2 down	up at 29, 11	up at 10	step
	14-5	1.75	2.6 up .9 down	2.1 (+) 1.3 (-)	9.7 (+) -6.4 (-)	412 up 407 down	18 up 15.2 down	up at 5	up at 5	step
	12-5	2.25	3.4 up 1.1 down	2.8 (+) 1.6 (-)	9.2 (+) -6 (-)	433 up 419 down	14.6 up 11.6 down	up at 5	up at 5	step
	7-1	9.5	14.3 up 4.8 down	11.7 (+) 6.9 (-)	10.7 (+) -6.8 (-)	406 up 414 down	4.9 up 3.7 down	up at 5	n/a	n/a
2 Contentious Activities	31-18	3	4.5 up 1.5 down	3.7 (+) 2.2 (-)	9.6 (+) -6.6 (-)	400 up 407 down	11.9 up 9.6 down	n/a	down at 18	step
	23-11	1.75	2.6 up .9 down	2.1 (+) 1.3 (-)	9.7 (+) -6.4 (-)	412 up 407 down	18.6 up 15.2 down	up at 17,13,11	up at 11	step
	17-1	3.25	4,9 up 1.6 down	4 (+) 2.7 (-)	11 (+) -6 (-)	569 up 403 down	11.9 up 8.7 down	up at 4	n/a	n/a
3 Violence Toward	31-10	0.5	.8 up .3 down	.6 (+) .4 (-)	7.4 (+) -6.2 (-)	400 up 417 down	33.1 up 49.7 down	up at 22,13,10	up at 10	step
	22-6	1.75	2.6 up .9 down	2.1 (+) 1.3 (-)	9.7 (+) -6.4 (-)	412 up 407 down	18.6 up 15.2 down	up at 10, 6	up at 6	step
	13-1	2.25	3.4 up 1.1 down	2.8 (+) 1.6 (-)	9.2 (+) -6 (-)	433 up 419 down	14.6 up 11.6 down	up at 4	n/a	n/a

Table 6. Consolidated Individual Univariate Analysis on Reversed SFOR Incident Data. Up corresponds to upward shifts and down corresponds to downward shifts.

Conducting the multivariate analysis of the reversed SFOR data is consolidated below in Table 7. As with the multivariate analysis on the SFOR data in its original order, there were no multivariate shifts detected.

			INDIVIDUAL UNIVARIATE PARAMETERS							
Data Category	Time Periods	Target In Control Mean	Out of Control Mean	k+/k-	h+/h- (DI)	In Control ARL	Out of Control ARL	Isolated Departures	Persistent Shifts	Type of Persistent Shift
1	31-8	1.5	2.3 up .8 down	1.9 (+) 1.1 (-)	11 (+) -8.1 (-)	1644 up 1644 down	25.4 up 25.1 down	up at 29, 11	up at 8	step
2	31-8	3	4.5 up 1.5 down	3.75 (+) 2.2 (-)	12.5 (+) -8.8 (-)	1657 up 1761 down	16.2 up 12.7 down	up at 13,11	n/a	n/a
3	31-8	0.5	.8 up .3 down	.6 (+) .4 (-)	11 (+) -9 (-)	1677 up 1743 down	51.1 up 77.4 down	up at 22,13,10	n/a	n/a
1	13-5	2	3 up 1 down	2.5 (+) 1.4 (-)	12.5 (+) -7.6 (-)	1973 up 1826 down	23.1 up 17.7 down	up at 5	up at 5	step
2	13-5	7	10.5 up 3.5 down	8.6 (+) 5 (-)	14.2 (+) -9 (-)	1646 up 1985 down	8.1 up 6.3 down	n/a	n/a	n/a
3	13-5	2.25	3.4 up 1.1 down	2.8 (+) 1.6 (-)	12.4 (+) -8 (-)	1733 up 1852 down	19.9 up 15.6 down	n/a	n/a	n/a
1	7-1	9.5	14.3 up 4.8 down	11.75 (+) 6.9 (-)	14 (+) -8.8 (-)	1640 up 1807 down	6.2 up 4.9 down	up at 5	n/a	n/a
2	7-1	7	10.5 up 3.5 down	8.6 (+) 5 (-)	14.2 (+) -9 (-)	1646 up 1985 down	8.1 up 6.3 down	up at 4	n/a	n/a
3	7-1	4.25	6.4 up 2.1 down	5.25 (+) 3 (-)	13.5 (+) -9 (-)	1692 up 1747 down	12 up 9.9 down	n/a	n/a	n/a

			NONPARAMETRIC MULTIVARIATE PARAMETERS						
Periods	k+	k-	Confidence Interval	Winsorizing Constant	Iterations	Start Point	Isolated Departures	Persistent Shifts	Type of Persistent Shift
31-8	3.75	1	99.94%	10	4800	7	n/a	n/a	n/a
13-5	3.75	1	99.94%	10	4800	7	n/a	n/a	n/a
7-1	3.75	1	99.94%	10	4800	7	n/a	n/a	n/a

Table 7. Consolidated Multivariate Analysis on Reversed SFOR Incident Data. Up corresponds to upward shifts and down corresponds to downward shifts.

The shifts that occurred during the multivariate analysis of the SFOR data in reverse order are shown below in Figures 26, 27, and 28.

These figures consolidate all the shifts that occurred during the multivariate analysis on one graph per data category.

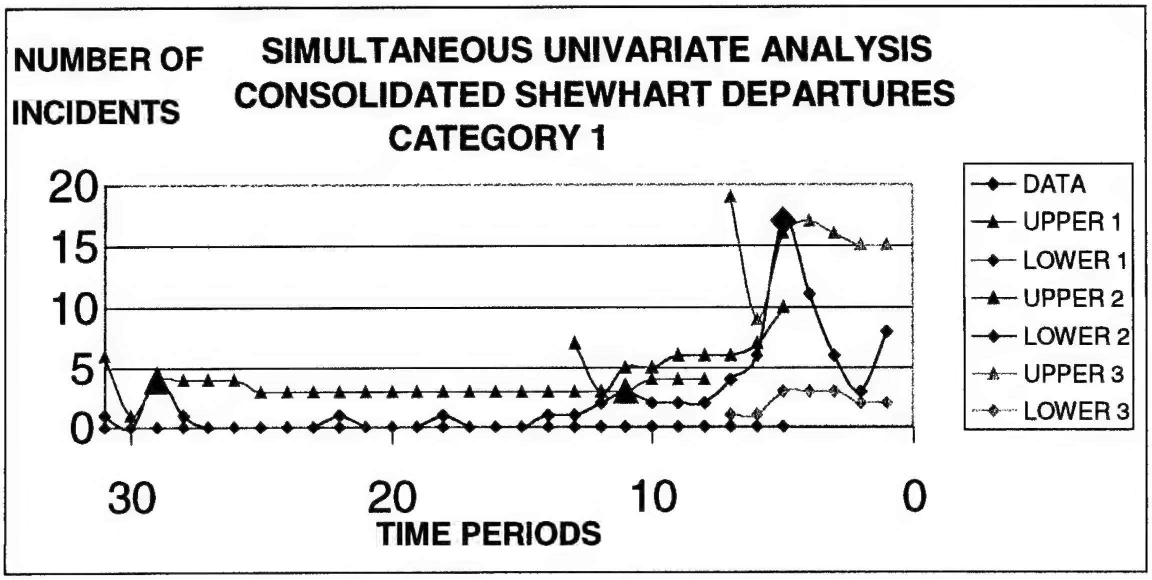

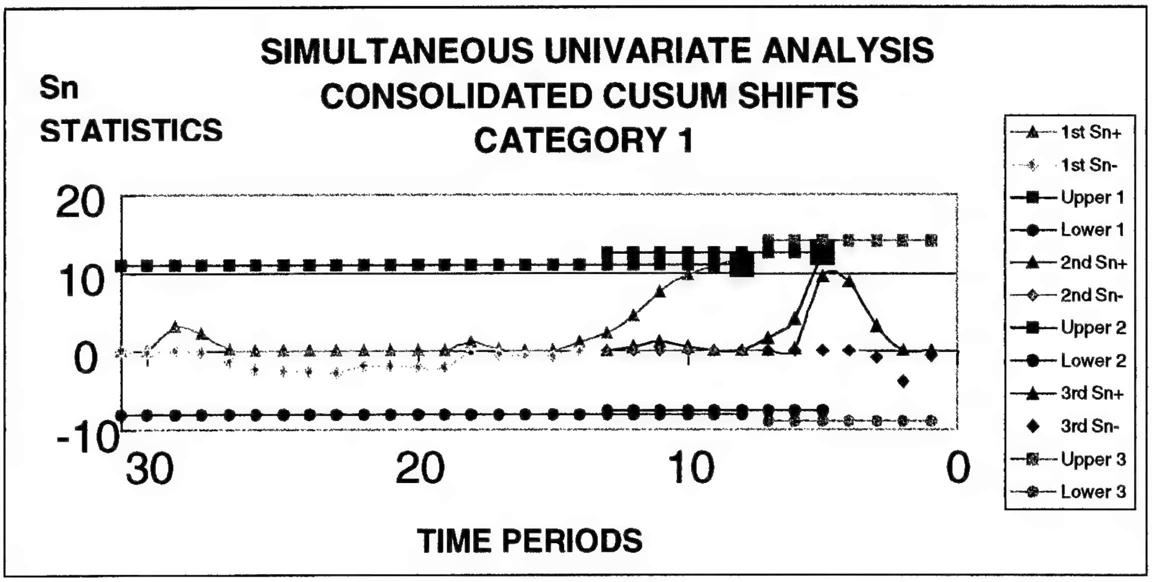

Figure 26. Multivariate Analysis for SFOR Data in Reverse Order, Consolidated Shifts in Category 1. 1st chart periods are from time period 31 to time period 8. 2nd chart periods are from time period 13 to time period 5. 3rd chart periods are from time period 7 to time period 1. Isolated shifts occurred in time periods 29, 11 and 5. Persistent shifts occurred in time periods 8 and 5. Large red data points identify departures and shifts.

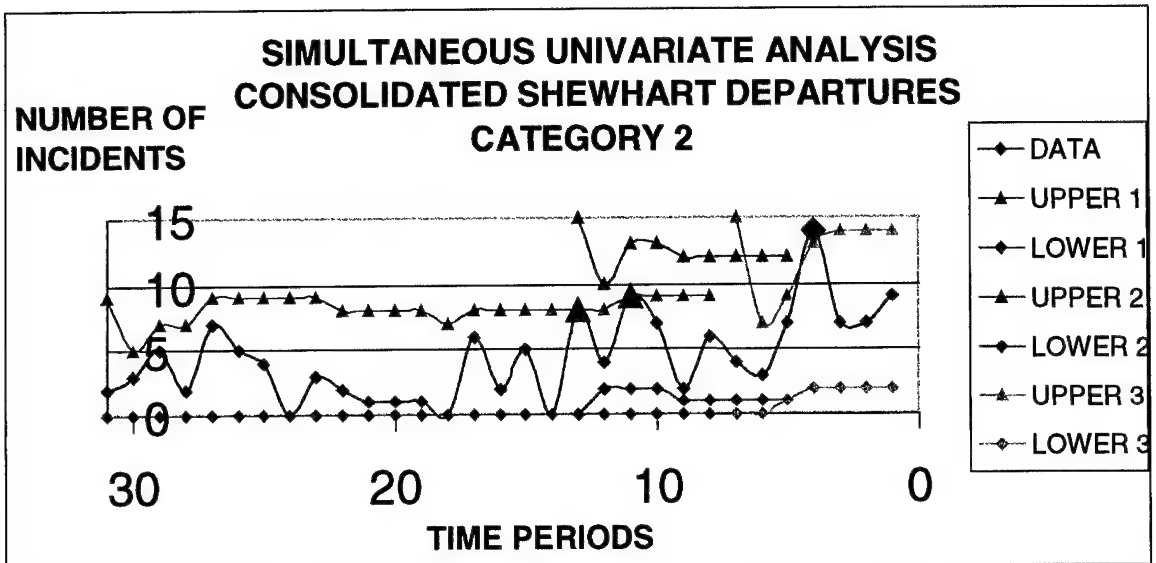

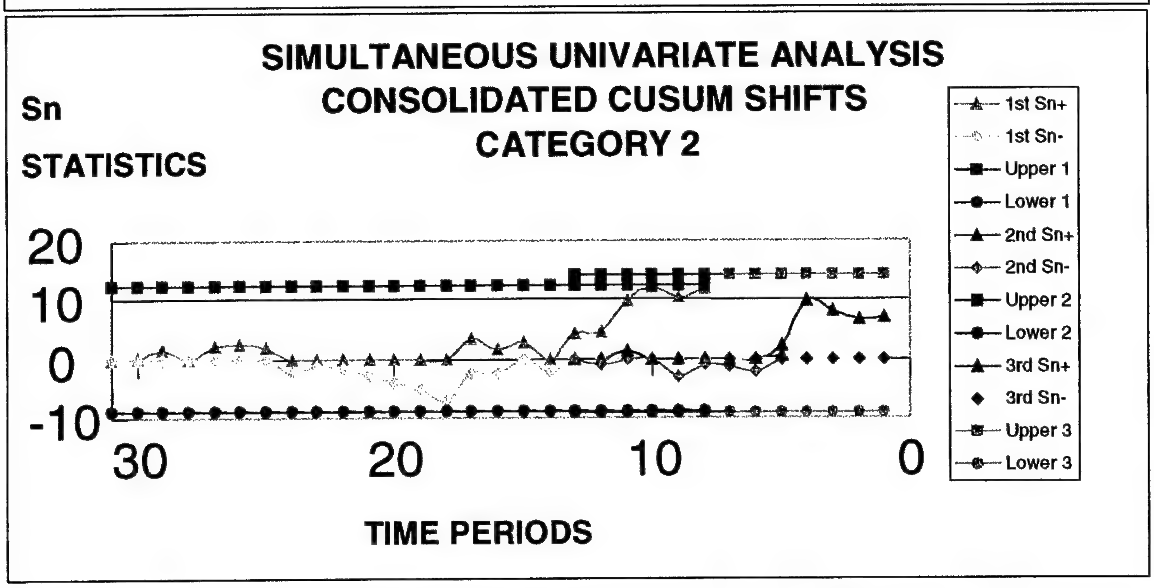

Figure 27. Simultaneous Univariate Analysis for SFOR Data in Reverse Order, Consolidated Shifts in Category 2. 1st chart periods are from time period 31 to time period 8. 2nd chart periods are from time period 13 to time period 5. 3rd chart periods are from time period 7 to time period 1. Isolated shifts occurred in time periods 13, 11 and 4. No persistent shifts were detected. Large red data points identify departures and shifts.

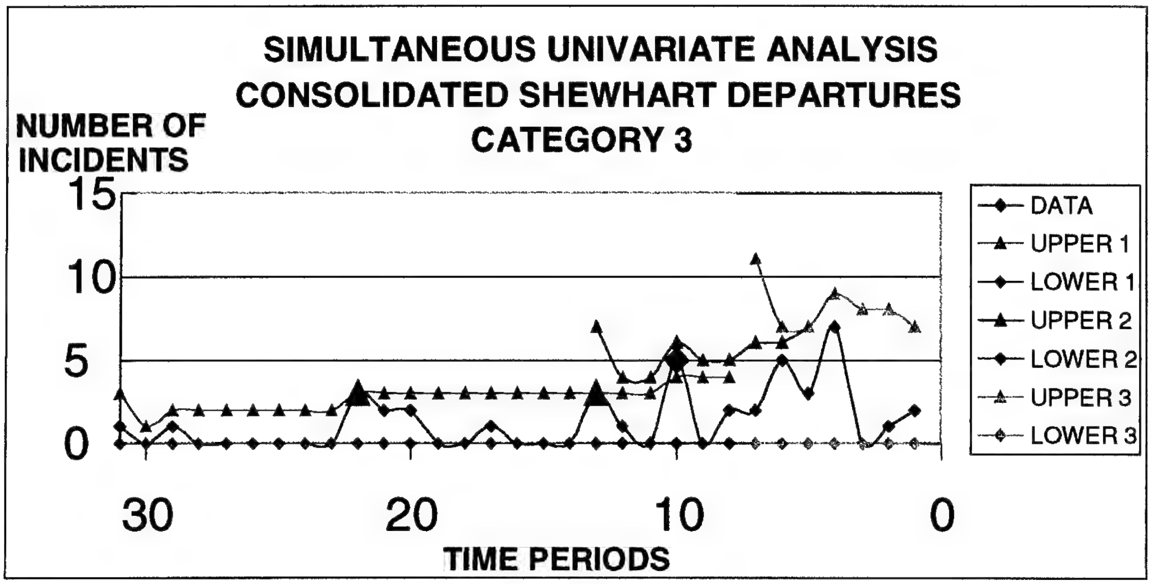

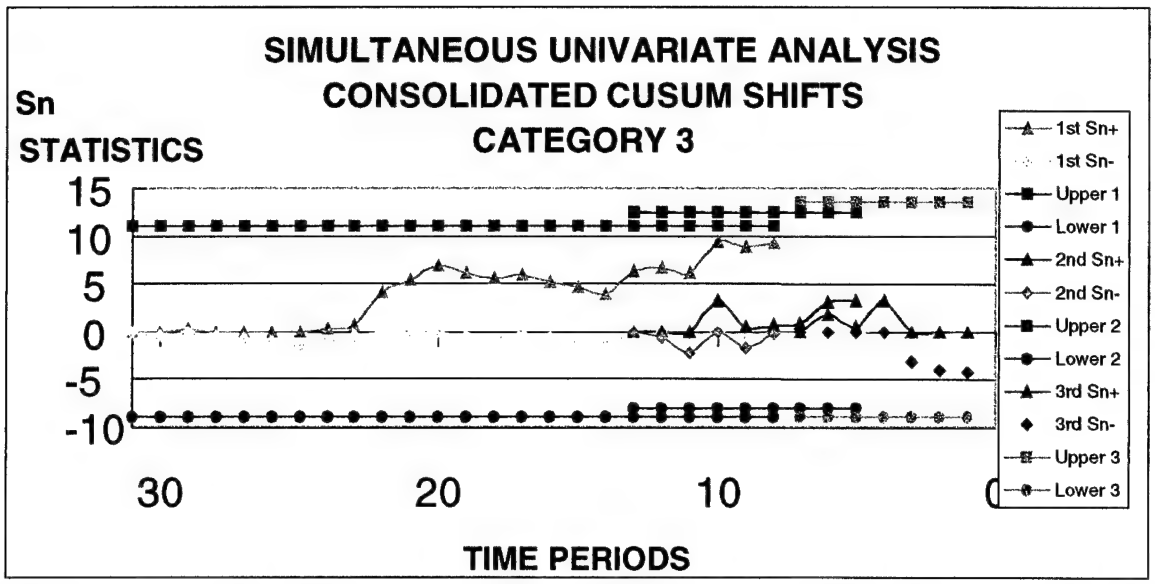

Figure 28. Simultaneous Univariate Analysis for SFOR Data in Reverse Order, Consolidated Shifts in Category 3. 1st chart periods are from time period 31 to time period 8. 2nd chart periods are from time period 13 to time period 5. 3rd chart periods are from time period 7 to time period 1. Isolated shifts occurred in time periods 22, 13 and 10. No persistent shifts were detected. Large red data points identify departures and shifts.

As could be expected, the general trends in the reversed data are similar but in the opposite direction of those in the actual data. In the individual univariate analysis, the three categories had a total of four persistent shifts, all of which downward shifts. In the reversed data, there were seven persistent shifts, six upward and one downward.

The difference in the number of shifts, the time periods when the shifts were detected, and the time periods when the shifts appeared to start can be explained by the different orderings of the data when reversed and its effects on the charts. Reversing the ordering of the SFOR data results in different time periods being used to calculate the initial target in control means and target out of control means. These will in turn result in slightly different upper and lower control limits, ARL's, and values of the calculated cumulative statistics. Combining the different ordering of the data with slightly different control limits will result different shifts on the control charts.

In the multivariate analysis, both the reversed and the actual data had two simultaneous univariate persistent shifts that necessitated the charts being retuned and restarted. Again the shifts were in opposite directions for the two data sets. The shifts in the actual data were all downward shifts; where as the shifts in the reversed data were all upward shifts.

The exercise of reversing the data is enlightening in that it clearly shows that the charts are effective in identifying upward shifts in the data, which for the SFOR commander in Bosnia has more significance and costly consequences than identifying downward shifts.

4. Conclusions on Analysis of SFOR Incident Data

Results from the analysis suggest several key issues about the situation that the commander should find informative and useful when developing his force protection plan. First, the situation was the most hostile in the initial data collection periods, 1 March through 5 April 1999, as denoted by high number of incidents in all data categories. The high numbers of enemy incidents were not naturally occurring random variations in the situation, but were instead statistically significant isolated departures from the normally observed values as shown by the departures signaled on the Shewhart charts. In particular, isolated upward departures in both the individual univariate and simultaneous

univariate Shewhart control charts occurred in category 3, violence towards SFOR, during time period 4, and in category 3, threats and rhetoric, during time period 5. Initial analysis for the possible causes of these incidents revealed that these isolated departures coincide with the United Nation's efforts to broker a peace settlement in Kosovo from February through the middle of March 1999, and the NATO air strikes against Serbian facilities, which commenced on 25 March 1999. These actions are likely to generate a negative responses from ethnic Serbians living in Bosnia. This negative response can be seen by looking at the SFOR incident log during 22 through 28 March, which corresponds to the start of the bombing campaign. The data log reveals that at least six of the eleven demonstrations against SFOR were anti-bombing demonstrations. From 29 March through 4 April, the number increased to 12 out of 17.

The high levels of enemy incidents explained above were isolated occurrences, with the numbers of incidents decreasing rapidly after 5 April. Increasing force protection levels after these incidents occurred is somewhat ineffective. The changes would not take effect until after the highest threat has already passed. If the increases in force protection were implemented in time period 5, they would be ineffective against the isolated upward departure in violence towards SFOR that occurred during time period 4. The increase in force protection levels would be effective in protecting the force against the decreasing but still high threat that was present from time period 5 through time period 8, 29 May through 25 April.

Commanders should not be completely convinced by this seemingly obvious cause of the high number of incidents. They should proceed with additional analysis of the situation to determine if other factors were present that may have caused or assisted in the increased number of incidents. The commander should use these factors to predict future enemy threat levels in similar situations. From these predictions,

commanders can initiate the appropriate force protection levels prior to the situation occurring, thus better protecting his unit. For example, if the commander knew in advance of another large bombing campaign against Serbian facilities in Serbia or Kosovo, he could increase the force protection levels based off of the number of incidents observed during time periods 4 and 5. This will at least give the commander an approximation to the possible threat level he will face in response to the new bombing campaign.

Secondly, the initial high hostility periods were followed by a continual decrease in the number of enemy incidents in all data categories through the end of the data collection period, 3 October 1999. The number of incidents decreased rapidly during time periods 6, 7, and 8. After time period 8, 25 April, the numbers of incidents appeared to stabilize. The tool developed in this thesis however, identified numerous statistically significant persistent decreases in the number of incidents after 25 April.

Both the individual univariate analysis and the simultaneous univariate analysis signaled persistent downward shifts in all data categories after time period 8. Individual univariate analysis identified the first persistent downward shifts as starting in time periods 6, 14, and 11, for the three data categories respectively. An additional persistent downward shift occurred in category 1, and appeared to start at time period 7. Simultaneous univariate analysis detected two persistent downward shifts in the three data categories. The first shift was detected in category 1, threats and rhetoric, at time period 9. The second persistent downward shift occurred in category 2, contentious activities, at time period 21. These shifts appeared to start in time periods 5 and 13 respectively.

All of these persistent decreases justify lowering the force protection level of the unit. The commanders and their staffs need to analyze the situation further to determine the specific causes of these

decreases and the appropriate force protection levels. By identifying the possible causes of these decreases, commanders could also focus their peacekeeping efforts in order to continue these trends.

It should be noted that there were two isolated departures detected following time period 8. The first was a downward departure in category 2, contentious activities, at time period 14, and the second was an upward departure in category 1, threats and rhetoric, during time period 29. As with other isolated departures discussed earlier, the causes of these departures should be determined and used for future reference.

Finally, the correlation between the data categories did not change. The fact that the nonparametric multivariate analysis did not detect any shifts in the correlation of the data categories suggests that the enemy's efforts, as divided among the three categories, remained constant. It can also be seen by the simultaneous increasing or decreasing trends that occurred in all three data categories. If a shift in the correlation between the data categories was detected, it would indicate a change in the enemy's distribution of effort. If the shift, for example, was from threats and rhetoric to acts of violence, the impact on force protection level would be significant. Identifying changes in the correlation is crucial to the commander in his assessment of the threat and his determination of appropriate force protection levels.

It is certain from the number of departures and shifts detected that the situation is volatile. The magnitude of this volatility is not realized, however, without comparing the number of shifts detected to the desired ARL's of the charts. The desired combined ARL's, or target false alarm rate, were 100 for each type of analysis. From this, one would expect one false alarm signal per independent univariate analysis data category and one false alarm signal in all multivariate analysis charts in 100 time periods or just over 2 years. Multiple shifts

occurred in both independent univariate analysis and multivariate analysis in only 31 time periods. This equates to a shift detection rate that is 3 to 6 times higher than the expected false alarm rate, depending on the data category. This amount of volatility is considerably larger than one might expect from just looking at the data. The tool developed in this thesis clearly identifies this high volatility in the SFOR data set. The commander must be made aware of such volatility if he is to make the initiate the proper force protection levels.

The overall recommendation after analyzing the SFOR incident data is that the force protection measures be reduced due to the statistically significant persistent decreases in the number of enemy incidents after 5 April 1999, time period 8. However, sufficient protection should be maintained to safeguard against possible isolated increases in enemy incidents, as detected in category 1, threats and rhetoric, during time period 29. Also, in the event that a similar bombing campaign is started against Serbian facilities, the commander should increase force protection levels based off the levels of enemy incidents seen previously, as in time periods 4 through 8.

THIS PAGE INTENTIONALLY LEFT BLANK

VI. CONCLUSIONS AND RECOMMENDATIONS

A. CONCLUSIONS

The methods and techniques developed and applied in this thesis, both the univariate SPC methods and the multivariate nonparametric permutation technique, effectively identified statistically significant changes in OOTW environments that might not have detected by current analysis methods. Current analysis methods are based on pattern recognition of enemy actions when compared to their doctrine. This is difficult in OOTW environments where enemy doctrine is often lacking if it exists at all. Pattern recognition methods do not differentiate between random fluctuations in the situation and statistically significant changes in the situation. This analysis is left to the commander who must rely on intuition and experience to determine if a significant change has occurred and the appropriate response to the change.

The use of SPC and the nonparametric multivariate technique developed in this thesis in the analysis of enemy incident data widens the applicability of SPC methods to an area of vital concern to the military, force protection. The effective application of these techniques not only provides commanders with the type of change that occurred in the situation, but also identifies the likely time at which the change started. From this information, the commander can focus his standard intelligence analysis to determine the causes of the shift, which can be used as the basis of his future plans and force protection levels. The information gained when using this analysis tool will be indispensable to commanders and staffs who are charged with conducting difficult missions in hazardous environments, while maintaining the security and safety of their soldiers.

This thesis combined standard univariate SPC analysis methods along with a technique for the nonparametric analysis of multivariate

data into a single statistical tool called "Multivariate CUSUM". Multivariate CUSUM was created with "ease of use" in mind. This was done to allow staff officers with basic training in statistics and SPC to manage the analysis of incident data and brief the results to their commander. Although the theory may be too complex for the untrained staff officer, trained personnel from the higher command levels will be able to educate their subordinate staff officers on the operation and application of Multivariate CUSUM, especially the graphical output. Once this is accomplished, the trained personnel will be able to monitor and supervise the subordinate staff's application of Multivariate CUSUM with minimal effort.

Multivariate CUSUM is implemented in Microsoft Excel, which is compatible with Army computer systems down to battalion level. It can easily be loaded on current Army computer systems and can be deployed with the unit wherever it may go.

Multivariate CUSUM is the first statistical tool to be offered for the analysis of the enemy situation in OOTW. It can augment current analysis methods to ensure the commander get the most complete and comprehensive estimate of the enemy situation possible. This tool and the information it provides will enable commander to make the appropriate and timely force protection decisions to ensure the safety and security of his soldiers.

B. RECOMMENDATIONS

As the number of Army OOTW missions continue to increase, the importance of force protection for deployed soldiers becomes more important. The IPB process alone is not sufficient to meet this challenge. Commanders need additional tools to assist them in determining the correct force protection posture for their unit. Multivariate CUSUM is a first step in meeting this challenge and ensuring the preparedness of deployed units and the safety of our soldiers.

Multivariate CUSUM is not a cure-all. It does not replace the need for the commander to know the abilities of his unit and the threats faced in the current situation when determining the force protection posture of his unit. Multivariate CUSUM is effective in identifying statistically significant changes in the current situation, which will improve the ability of the commander to properly assess the best force protection level for his unit and to better protect his soldiers.

Multivariate CUSUM should be fielded and deployed with the higher headquarters of deploying units, division and above, in sufficient time for the personnel and the commander to become trained on its use. Sufficient time must also be allowed for the controlling staff to brief their subordinates on its use since the subordinate units will be the units gathering the data. Without consistent and proper data collection, any analysis will be questionable.

C. TOPICS FOR FURTHER STUDY

Additional study could be conducted to determine an efficient method of calculating the Out of Control ARL's for multiple possible shifts in the multivariate analysis. This would give the commanders better insight into the time required for the technique to signal a given target shift in the data and assist in power calculations.

Also, further research could be conducted to develop a method to assist in determining the values of k^+, k^-, and the Winsorizing constant. Simulation was used in this thesis to identify acceptable values for these parameters. A statistical or mathematical method would be more efficient and give the user a more deterministic means of calculating the parameters.

Multivariate CUSUM is designed for the analysis of three variables. Additional work could be done to scale the program for analysis of an arbitrary number of variables.

Finally, further research could be done to determine the applicability of these methods into the area of friendly unit deception.

If the enemy were to use a similar tool, he may be able to make more precise predictions on our future actions and therefore better prepare to defeat them. Multivariate CUSUM may be effective in identifying the predictability of our actions and deception plans. By self-analyzing our actions and plans, we may prevent the enemy from identifying changes in our posture and preparing against our future actions.

APPENDIX A. SFOR INCIDENT LOG SUMMARY

This data was taken from March through October 1999 from the SFOR incident log at Task Force Eagle. Entries into the log that did not pertain to local populace actions toward SFOR units were disregarded.

CONSOLIDATED SFOR INCIDENT DATA

March - October 1999

Month	Dates	Time Periods	Category 1 Threats & Rhetoric	Category 2 Contentious Activities	Category 3 Violence Towards SFOR
March	1-7	1	8	9	2
	8-14	2	3	7	1
	15-21	3	6	7	0
	22-28	4	11	14	7
April	29-4	5	17	7	3
	5-11	6	6	3	5
	12-18	7	4	4	2
	19-25	8	2	6	2
May	26-2	9	2	2	0
	3-9	10	2	7	5
	10-16	11	3	9	0
	17-23	12	2	4	1
	24-30	13	1	8	3
June	31-6	14	1	0	0
	7-13	15	0	5	0
	14-20	16	0	2	0
	21-27	17	0	6	1
July	28-4	18	1	0	0
	5-11	19	0	1	0
	12-18	20	0	1	2
	19-25	21	0	1	2
	26-1	22	1	2	3
August	2-8	23	0	3	0
	9-15	24	0	0	0
	16-22	25	0	4	0
	23-29	26	0	5	0
September	30-6	27	0	7	0
	6-12	28	1	2	0
	13-19	29	4	5	1
	20-26	30	0	3	0
October	27-3	31	1	2	1

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX B. DIRECTIONS FOR USING MULTIVARIATE CUSUM

A. GENERAL

1. Begin any analysis by entering the data into the "data1" page. Column A is a number that designates the time period. Columns B, C, and D are the actual data values of the time period.

2. When restarting the charts and updating the time periods and data values on "data1" page, use only the "Paste Special Values" option in Excel.

B. UNIVARIATE ANALYSIS

1. Press the "F9" key to calculate the target in control lambda's, the lambda+s', and the lambda-'s for the different data categories. These values are currently set for a 50% increase and a 50% decrease of the target in control lambda for each category. This targeted shift may be change at the user's discretion by changing the underlying equations in the appropriate cells.

2. Press the "Run GETH" command button to execute ANYGETH.exe and determine the CUSUM chart control limits. Directions for using ANYGETH.exe are in Appendix C.

3. Press the appropriate "Change Parameters _" command button for each of the data categories. Enter the decision intervals obtained from ANYGETH.exe into the Upper limit and Lower limit windows. Enter the target Lambda in control, the Lambda+, and the Lambda- from the appropriate cells on the Excel "data1" page for the corresponding data category. Enter the desired Shewhart chart probability limit into the Isolate Probability Limits window. Press the "OK" command button when complete.

4. Select the "Update Univariate Graphs" command button to update the univariate graphs. Multivariate will take you to the univariate graphs of data category 3. You can move to the other graphs by selecting the appropriate worksheet tab at the bottom of the Excel window or move back to the "data1" page by selecting the "Go to Data" command button.

5. If a category goes out of control, the charts will plot the points outside the control limits. The "data1" page will also display the work "hot" in the appropriate time period for the corresponding data category. Charts do not have to be retuned and restarted for isolated shifts. They do have to be retuned and restarted for persistent shifts.

C. MULTIVARIATE ANALYSIS

1. Conduct simultaneous univariate analysis in the same manner describe above in univariate analysis being sure to start all charts when a persistent shift is detected in any one of the CUSUM charts.

2. Once the simultaneous univariate analysis is complete, return to the Excel "data1" page. Enter the desired values for the parameters listed beneath the "Update Multivariate Graphs" button.

a. Recommend starting with values of k+ and k- equal to 4 and 2 respectively. After executing the "Update Multivariate Graphs" command button below, the values of k+ and k- should be adjusted to obtain the appropriate control limits.

b. Recommend a Winsorizing constant equal 10. As with the values of k+ and k-, the Winsorizing constant should be adjusted after executing the "Update Multivariate Graphs" command button below to obtain the appropriate control limits.

c. Recommend an initial number of permutations equal to 500. The user will save time by running smaller number of permutations when adjusting the k+, k- and Winsorizing parameters. When these parameters are appropriate, this thesis recommends running 4800 permutations to obtain smooth control limits and thorough sampling of the data.

d. With 3 data categories, this thesis recommends an initial an initial starting point of 7. Although this does not totally remove problems caused by near-singular covariance matrices, it sufficiently reduces the problem without sacrificing data observations.

e. The confidence interval of the multivariate charts is based on the desired ARL. 99.94% was used in this thesis to achieve a multivariate test ARL of 1667 and an overall process ARL of 100.

3. Once the parameters are updated, select the "Update Multivariate Graphs" command button to begin the nonparametric permutation technique and to update the univariate graphs. Multivariate will take you to the multivariate Shewhart control chart. You can move to the other graphs by selecting the appropriate worksheet tab at the bottom of the Excel window or move back to the "data1" page by selecting the "Go to Data" command button.

4. If a category goes out of control, the charts will plot the points outside the control limits. The "data1" page will also display the work "hot" in the appropriate time period in the "Multivariate Hot" columns. Once again, charts do not have to be retuned and restarted for isolated shifts. They do have to be retuned and restarted for any persistent shifts, either from the univariate charts or from the multivariate charts.

5. When restarting the charts because of a persistent shift, all data categories are started at the same time regardless of whether or not they are out of control. Follow the steps listed above for univariate analysis and multivariate analysis to restart the charts and conduct analysis on the new time periods.

APPENDIX C. DIRECTIONS FOR USING ANYGETH.EXE

1. Open ANYGETH.exe from the Visual Basic command button.

2. Select the desired distribution from the provided list. For example, if the Poisson Distribution desired, enter the number 3 and press return.

3. Enter the desired target in control mean and out of control mean. Separate the values by a space or a carriage return. In Multivariate, the in control means are calculated on "data1" in cells J9 for Category 1, J15 for Category 2, J21 for Category 3. Target out of control means for an upward shift of 50% of the in control mean are calculated on "data1" in cells J10 for Category 1, J16 for Category 2, J22 for Category 3. Target out of control means for a downward shift of 50% of the in control mean are calculated on "data1" in cells J11 for Category 1, J17 for Category 2, J23 for Category 3.

4. ANYGETH.exe will calculate the exact theoretical reference value. This value should be rounded because the ANYGETH.exe may not converge on an appropriate decision interval using the exact theoretical reference value. Recommend rounding to the nearest 10^{th}. For example, if ANYGETH.exe returns a theoretical reference value of 4.23, round the number to 4.2 and press return.

5. Enter -999 999 to execute ANYGETH.exe without a Winsorzing Constant. For information regarding Winsorization in Statistical Process Control, refer to *Cumulative Sum Charts and Charting for Quality Improvement* by D. Hawkins and D. Olwell.

6. Select the desired chart, either "z" for zero start CUSUM or "f" for Fast Initial Response and press return. This thesis uses zero start CUSUM charts exclusively.

7. Enter the appropriate average run length (ARL) and press return. This thesis uses a test ARL of 1600 to obtain an overall process ARL of 100.

8. ANYGETH.exe will calculate the appropriate control limit. This value is designated as the Decision Interval. ANYGETH.exe always returns a positive Decision Interval value. The lower Decision Interval values should be entered as negative values when input into Multivariate. For example, ANYGETH.exe returns a lower Decision Interval of 4.4, the user should input -4.4 when entering the values into the Multivariate "Change Parameter" window.

9. Repeat the above steps for each upper and lower control limit for each data category.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX D. VERIFICATION OF POISSON DATA

The table below shows the results of two separate tests that attempt to verify that the data is from the Poisson distribution. The first test is the "mean equals variance" test. This general test for Poisson data tests if the mean of the sample is generally close to the variance of the sample. This follows from the property of Poisson data that the mean is equal to the variance. The test shows variances that are generally twice as large as the means of the samples. This would suggest that the data is not Poisson. However, this may be explained by the presence of multiple Poisson processes. If multiple Poisson processes are present, the variance will be larger than the mean of the sample. This is because the tails of the individual Poisson distributions will spread out the variance of the combined sample.

The second test is the χ^2 Goodness of Fit Test. This test is a more precise test than the "mean equals variance" test. The results of this test show that the data may be plausibly Poisson, as the *p* values obtained were larger than the alpha used for the test, 0.01. One limitation of this test when used on this data set is that it requires the data to have a constant mean. As shown in this thesis, the means of all data categories changed throughout the 31 time periods. This resulted in the 31 sample periods being reduced to generally the largest in control sample of each variable. For example, data category 1 had the longest run in control from time period 13 to time period 31 as shown by the box around the data. This was the sample size used for the test.

Another weakness of this test when used on this data set is that the test requires bin sizes larger than 5. Dividing the small in control sample sizes into three bins, resulted in numerous bin sizes that were close to or equal to 5.

Given the limitations and discrepancies of these two tests, this thesis concluded that the data may be plausibly Poisson.

Period	Cat 1	Cat 2	Cat 3
1	8	9	2
2	3	7	1
3	6	7	0
4	11	14	7
5	17	7	3
6	6	3	5
7	4	4	2
8	2	6	2
9	2	2	0
10	2	7	5
11	3	9	0
12	2	4	1
13	1	8	3
14	1	0	0
15	0	5	0
16	0	2	0
17	0	6	1
18	1	0	0
19	0	1	0
20	0	1	2
21	0	1	2
22	1	2	3
23	0	3	0
24	0	0	0
25	0	4	0
26	0	5	0
27	0	7	0
28	1	2	0
29	4	5	1
30	0	3	0
31	1	2	1

Time Periods 8-31		
Mean 1	Mean 2	Mean 3
0.52631579	2.88235294	0.875
Variance 1	Variance 2	Variance 3
0.92982456	4.48529412	1.76630435
n-l = 18	n-1 = 16	n-1 = 24

	CHI 2 GOF	CHI 2 GOF	CHI 2 GOF	Chi 2 stat	alpha = .01
	0.4585	0.8932	4.778	6.6348913	df = 1
p value	CHI SQRD	CHI SQRD	CHI SQRD	alpha	
	0.49832584	0.34461162	0.02882558	0.01	
	PLAUSIBLY POISSON				
	yes	yes	yes		
	fail to reject	fail to reject	fail to reject		

POISSON IF:
CHI 2 GOF < CHI 2 stat
or
CHI SQRD > alpha

LIST OF REFERENCES

Air Force News. 1996. *DOD Announces New Military Force Protection Measures.* Available at Air Force Link web site: http://www.af.mil/news/Sep1996/n19960919_960951.html. Internet. Accessed 24 October 1999.

Anderson, T.W. 1971. *An Introduction to Multivariate Statistical Analysis.* 2d ed. New York: John Wiley and Sons.

Center for Army Lessons Learned (CALL). 1996. *Mountain Eagle Lessons: Stability operations by the Combat Maneuver Training Center (CMTC), 4th QTR FY95 &1st QTR FY96* [Report on line]. Available from http://www.call.army.mil/call/ctc_bull/llcmtc96/toc.htm. Internet. Accessed 15 October 1999.

Cohen, William S. 1997. *Report: Personal Accountability for Force Protection at Khobar Towers* [Report on line]. Available from Defense Link web site: http://www.defenselink.mil/pubs/khobar/report.html. Internet. Accessed 24 October 1999.

Cook, R. Dennis and Sanford Weisberg. 1999. *Applied Regression Including Computing and Graphics.* New York: John Wiley and Sons.

Crosier, Ronald B. 1988. Multivariate Generalizations of Cumulative Sum Quality-Control Schemes. *Tecnometrics.* Vol. 30. No 3.

Department of the Army. 1990. *FM 34-3, Intelligence Analysis.* Washington, D.C.: Government Printing Office.

Department of the Army. 1993. *FM 34-7, Intelligence and Electronic Warfare Support to Low-Intensity Conflict Operations.* Washington, D.C.: Government Printing Office.

Department of the Army. 1994. *FM 100-23, Peace Operations.* Washington, D.C.: Government Printing Office.

Devore, Jay L. 1995. *Probability and Statistics for Engineering and the Sciences.* Pacific Grove: Brooks/Cole.

Hawkins, Douglas M. and David H. Olwell. 1998. *Cumulative Sum Charts and Charting for Quality Improvement.* New York: Springer.

Johnson, R. and D. Wichern. 1992. *Applied Multivariate Statistical Analysis.* New York: Prentice Hall.

Law, Averill M. and W. David Kelton. 1991. *Simulation Modeling & Analysis.* 2d ed. New York: McGraw-Hill Inc.

Montgomery, Douglas C. 1991. *Introduction to Statistical Quality Control.* 2d ed. New York: Wiley.

Moustakides, G.V. 1986. Optimal stopping times for detecting changes in distributions. *Annals of Statistics.* Vol. 14.

Olwell, David H. 1997. Managing Misconduct: Statistical Process Control Applied to Sexual Harassment. *1997 Proceedings of the Section on Quality and Productivity*. Alexandria, VA: American Statistical Association.

Perry, William J. 15 September 1996. *Memorandum for the Chairman joint Chiefs of Staff* [Memorandum on line]. Available from Federation of American Scientists web site: http://www.fas.org/irp/threat/downing/annx_d.html. Internet. Accessed 24 October 1999.

Ross, Sheldon M. 1997. *Introduction to Probability Models*. San Diego: Academic Press.

Schneider, David I. 1998. *An Introduction to Programming Using Visual Basic 5.0*. 3d ed. Upper Saddle River: Prentice Hall Inc.

Van Dobben de Bruyn, C.S. 1968. *Cumulative Sum Tests: Theory and Practice*. London: Griffin.

Weitzman, Robert C. 1999. Statistical Monitoring of Police Force For Rapid Detection of Changes in Frequency. Master's Thesis, Naval Postgraduate School, Monterey, California.

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center..2
   8725 John J. Kingman Rd., Suite 0944
   Fort Belvoir, VA 22060-6218

2. Dudley Knox Library..2
   Naval Postgraduate School
   411 Dyer Rd.
   Monterey, California 93943-5101

3. LTC David Olwell..1
   Operations Research Department (OR/OL)
   Naval Postgraduate School
   1411 Cunningham Road
   Monterey, California 93943

4. Professor Tom Lucas..1
   Operations Research Department (OR/OL)
   Naval Postgraduate School
   1411 Cunningham Road
   Monterey, California 93943

5. Mr. Walter Hollis..1
   DUSA (OR), HQ Dept of the Army
   Room 2E660 PNT
   102 Army Pentagon
   Washington DC 20350

6. Dr. Henry Dubin..1
   Assistant Secretary of the Army
   (Research, Development & Acquisition)
   Director, Assissment and Evaluation
   103 Pentagon
   Washington, D.C. 20310-0103

7. Chief, Strategic Plans..1
   MND(N), Task Force Eagle
   Tuzla, Bosnia, APO AE 09789

8. Mr. Michael Bauman..1
   Director, TRADOC Analysis Center
   Ft. Leavenworth, KS, 66027

9. Director, Center for Army Analysis 1
   Ft. Belvoir, VA 22060

10. Director, Center for Army Lessons Learned 1
    Ft. Leavenworth, Kansas 66043

11. LTC Tom Kastner..1
    7535 Will Street
    Ft. Meade, MD 20755

12. Dr. Doug Hawkins..1
    Chairman, Department of Applied Statistics
    University of Minnesota
    1994 Buford Avenue, (352 COB)
    St. Paul Minnesota 55108-6042

13. MAJ Paul J. Finken...1
101 3rd Avenue
Earling, Iowa 51530